

IntelligenceCenter User Guide

Version 3.3.1.1

BLUE COAT

Copyright/Trademarks/Patents/Disclaimer

© 2014 Blue Coat Systems, Inc. All rights reserved. BLUE COAT, PROXYSG, PACKETSHAPER, CACHEFLOW, INTELLIGENCECENTER, CACHEOS, CACHEPULSE, CROSSBEAM, K9, DRTR, MACH5, PACKETWISE, POLICYCENTER, PROXYAV, PROXYCLIENT, SGOS, WEBPULSE, SOLERA NETWORKS, DEEPSEE, DS APPLIANCE, SEE EVERYTHING. KNOW EVERYTHING., SECURITY EMPOWERS BUSINESS, BLUETOUGH, the Blue Coat shield, K9, and Solera Networks logos and other Blue Coat logos are registered trademarks or trademarks of Blue Coat Systems, Inc. or its affiliates in the U.S. and certain other countries. This list may not be complete, and the absence of a trademark from this list does not mean it is not a trademark of Blue Coat or that Blue Coat has stopped using the trademark. All other trademarks mentioned in this document owned by third parties are the property of their respective owners. This document is for informational purposes only.

BLUE COAT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT. BLUE COAT PRODUCTS, technical services, and any other technical data referenced in this document are subject to U.S. export control AND SANCTIONS LAWS, REGULATIONS AND REQUIREMENTS, AND MAY BE SUBJECT TO EXPORT OR IMPORT REGULATIONS IN OTHER COUNTRIES. YOU AGREE TO COMPLY STRICTLY WITH THESE LAWS, REGULATIONS AND REQUIREMENTS, AND ACKNOWLEDGE THAT YOU HAVE THE RESPONSIBILITY TO OBTAIN ANY LICENSES, PERMITS OR OTHER APPROVALS THAT MAY BE REQUIRED IN ORDER TO EXPORT, RE-EXPORT, TRANSFER IN COUNTRY OR IMPORT AFTER DELIVERY TO YOU.

America's:

Blue Coat Systems, Inc.

410 N. Mary Ave.

Sunnyvale, CA 94085

Rest of the World:

Blue Coat Systems International SARL

3a Route des Arsenaux

1700 Fribourg, Switzerland

Table of Contents

Getting Started with IntelligenceCenter	1
Set Up Data Collection and Reporting	2
Manage IntelligenceCenter	3
Log In to IntelligenceCenter	3
Log Out of IntelligenceCenter	5
Manage Audit Logging	5
Manage Licensing	7
Manage Scheduled Tasks	8
Configure System Settings	10
Troubleshoot IntelligenceCenter	20
Manage the Databases	21
Set Up Data Collection and Reporting	24
Data Collection Overview	24
Add DataCollector to the Network Group	26
Define Data Sources	28
Set Up Business Hours Reporting	32
Share Business Hours Settings Between Multiple DCs	33
Force Class Discovery	36
Modify Data Collection Settings	37
Back up FDR Data to a CSV File	43
Remove a Data Source	44
Monitor DataCollector Health	46
Manage User Access	51
User Access Overview	51
Add a User Profile	52
Modify a User Profile	54
Manage Your User Profile	55
Delete a User Profile	56
Configure User Groups	56
Enable External Authentication	59
Manage Roles	64
Manage Devices	71
IC Network Overview	71
Create a Group	73
Create a Sub-Group	73
Add a Device	75
Configure a Device	77

Import Devices from PolicyCenter	85
Configure Single Sign-On.....	87
Modify a Device Configuration	88
Remove a Device	90
Manage PacketShaper Configurations	91
Monitor Devices	94
View Device Information.....	94
Monitor DataCollector Health	96
View the PacketShaper Class Tree	101
Manage Network Views	102
Network Views Overview.....	102
Create a Network View.....	103
Modify a Network View	105
Remove a Network View	107
Manage Applications.....	108
Applications Overview	108
Create an Application	110
Define Critical Applications.....	113
Modify an Application	114
Remove an Application	115
Configure Static Reporting.....	116
Reporting Overview	116
Customize the Report List.....	120
Run a Report	121
Find a Report.....	131
View a Report	132
Delete a Report	133
Define Report Schedules	134
Export Report Data.....	134
Manage Scheduled Tasks.....	136
Manage Report Schedules.....	137
Create Consolidated Reports	143
Manage Sites	150
Add a Site	150
Remove a Site	151
Manage Alerting	152
Alert Type Configuration Details	152
View Alerts.....	153
Configure Global Alert Settings	155

Configure Alert Type Settings	157
Set Global Coalesce Interval.....	158
Set Global Deletion Interval.....	159
Set Global Notification by Severity.....	159
Send Alert Notifications to a Syslog Server	159
Send SNMP Traps	160
Enable or Disable Alerting.....	160
Configure Portlets	161
Manage Portlets	162
Manage Portlet Views	162
Add a Portlet.....	164
Configure a Portlet	164
View a Portlet	191
Remove a Portlet.....	191
Arrange Portlet Windows.....	192
Reports	193
IntelligenceCenter Reports	193
Application Reports	195
Application Activity Report.....	195
Application Response Time Report.....	197
Top Applications Report	202
Top Immediate Children Report	204
Device Reports.....	206
Device Compression Summary Report.....	206
Link Utilization Report	211
TCP Health Report	216
Top Traffic Classes Summary Report	220
Traffic Class Compression Report	222
Traffic Class Response Time Report	226
Traffic Class Utilization Report	230
VoIP Statistics Report	233
Host Reports	238
Host Pairs Activity Report.....	238
Top DSCP Report	240
Top Host Pairs Report.....	243
Top Listeners Report.....	245
Top Services Report.....	247
Top Talkers Report.....	251
Top VLAN Report	253

Site Reports	256
Site Response Time Report	256
Top Applications by Site Report	259
Top Host Pairs by Site Report	261
Top Listeners by Site Report	263
Top Services by Site	265
Top Sites Report	267
Top Talkers by Site Report	270
Drill-Down Reports	272
Drill Down: Application Activity	272
Drill Down: Application Rates Comparison	273
Drill Down: Host Pairs Activity by Application	274
Drill Down: Host Pairs Activity by Traffic Class	275
Drill Down: Host-Pair Flow Details	276
Drill Down: Top Talkers/Listeners for Class	278
Drill Down: Top Talker/Listener Peers	280
Drill Down: Top Talkers/Listeners for Application	282
Drill-Down: Peers for Listener/Talker	284
Drill-Down: Flow Details	285
Drill Down: Traffic Class Activity by Device or Group	287
Drill Down: Traffic Class Rates Comparison	288
Portlets	289
Application Performance Portlet	290
Class Utilization Portlet	294
Network Efficiency Portlet	297
Per Server FDR Portlet	299
Per Subnet FDR Portlet	301
Top N Children Portlet	304
VoIP Performance Portlet	306
Reference	309
Measurement Variables	309
Measurement Variables Collected from PacketShaper	309
Measurement Variables Collected from PacketShaper ISP	314
CSV Record Formats	316
Packeteer-2 Record Format	316
Report Options	320
Application Activity Report Options	321
Application Response Time Report Options	321
Device Compression Summary Report Options	321

Host Pairs Activity Report Options	322
Link Utilization Report Options	323
Site Response Time Report Options.....	324
Site Response Time Report Options.....	324
TCP Health Report Options.....	324
Top Applications Report Options.....	325
Top Applications by Site Report Options	326
Top DSCP Report Options	327
Top Host Pairs Report Options	328
Top Host Pairs by Site Report Options	329
Top Immediate Children Report Options.....	330
Top Listeners by Site Report Options	331
Top Listeners Report Options	332
Options for Top N by Site Reports	333
Top Services by Site Report Options	334
Top Services Report Options	335
Top Sites Report Options	335
Top Talkers by Site Report Options	336
Top Talkers Report Options	337
Top Traffic Classes Summary Report Options.....	338
Top VLAN Report Options.....	339
Traffic Class Compression Report Options.....	339
Traffic Class Response Time Report Options.....	339
Traffic Class Utilization Report Options	340
VoIP Statistics Report Options	341
Index	343

Getting Started with IntelligenceCenter

IntelligenceCenter (IC) is a central monitoring and reporting solution for Blue Coat products. IC provides a secure dashboard that allows access to Blue Coat products from a central screen. With a single click, you can go into PolicyCenter or PacketShaper—without having to enter a user name and password each time a product is accessed. Just log in to IC and the Blue Coat products are instantly accessible without additional sign on.

You can configure one or more DataCollectors (DCs) to collect [Measurement Engine \(ME or Metric\)](#) and/or [Flow Detail Record \(FDR\)](#) data from PacketShaper appliances on your network. You can also collect NetFlow-5 data from other network devices such as routers. You can then use the data to generate aggregated reports that provide snapshots of your network as a whole. Or, you can define your own applications and network views, which allow you to generate reports using specific cross-sections of data.

IC provides two types of reporting mechanisms:

- [Static Reports](#)—Allow you to choose from a selection of predefined reports, which you can run ad hoc or on a schedule.
- [Portlets](#)—Allow you to monitor and troubleshoot the flows and applications that are the most important to you. Portlet configurations are unique to each IC user profile.

There are several configuration steps that you must follow in order to get IntelligenceCenter up and running. To use IntelligenceCenter for reporting and/or single sign-on, you must:

1. [Set up data collection and reporting.](#)
2. [Create roles](#) that define the various access levels to the IC system.
3. [Create user and group accounts for your IC users.](#)
4. [Enable single sign-on.](#) In order to use the single sign-on capabilities, your end users must enter their sign-on credentials for the devices that they are authorized to access.

After you complete these steps, your users will be able to begin generating [reports](#) and [portlets](#) and using the single sign-on feature to access Blue Coat devices.

Set Up Data Collection and Reporting

Before you can start generating IntelligenceCenter (IC) reports, you must configure your IC network. At a minimum, an IC network contains a group of devices and a DataCollector (DC) that collects and reports on the data generated by these devices. Depending on the amount of data you need to collect and the way you want to partition your reporting, you may need to create multiple groups, each with its own DC and unique set of devices. Each DC collects data from and reports on the devices that reside in its own group.

To set up your IC network, you must complete the following tasks:

1. Define the [groups](#) and [sub-groups](#) that will contain the DCs and devices in your IC network. Each group you create represents a standalone reporting domain with a unique DC and set of devices.
2. [Add a DC](#) to each group.
3. [Add each device](#) that you want to be able to report on to the IC network.
4. [Configure DC to collect data from the devices](#).
5. (optional) [Define a set of business hours](#) to allow you generate reports that only include data for the times when the majority of network activity occurs.
6. [Define the applications that you are interested in reporting on and designate critical applications](#).
7. [Enable site reporting by associating IC site names with PacketShaper location-based traffic classes](#).

Manage IntelligenceCenter

Log In to IntelligenceCenter

You can log in to IntelligenceCenter from any client machine that is equipped with one of the following browsers:

- Google Chrome 35
- Firefox 30
- Internet Explorer 7.x, 8.x, or 9.x

In order for the IntelligenceCenter user interface to function properly, you must set the browser's Internet security settings to the default values.

In addition, because the IntelligenceCenter user interface is displayed using Adobe Flash Player, you must have the current version of Adobe Flash Player installed on the client system from which you will access IntelligenceCenter. If you haven't already installed the latest version, make sure to do so before using IntelligenceCenter. If you aren't sure which version of Adobe Flash Player is installed on your client system, go to:

<http://www.adobe.com/software/flash/about/>

To download the latest version, go to:

<http://www.adobe.com/products/flashplayer/>

To log in to IntelligenceCenter:

1. Open a browser window.
2. For the URL, type one of the following:

https://<ipaddress>:<port> (for a secure connection)

- or -

http://<ipaddress>:<port> (for a non-secure connection)

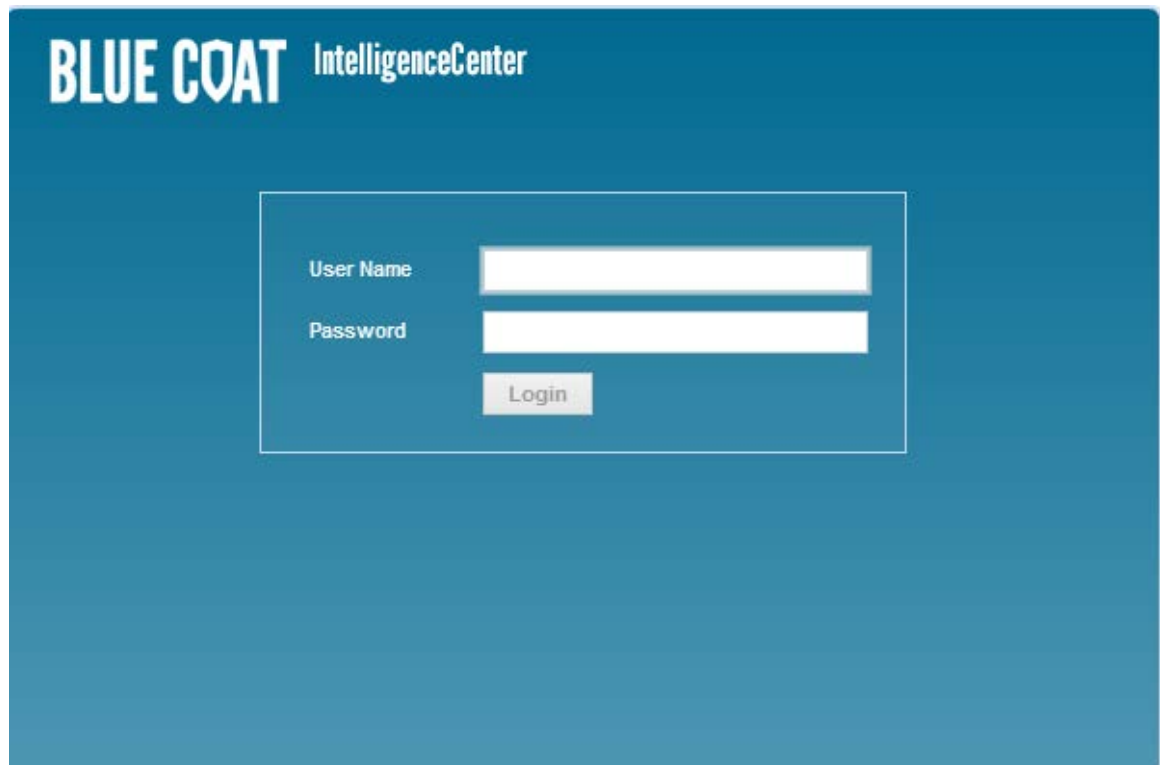
Note: The port number is required only if you specified HTTP and/or HTTPS port numbers other than the defaults — 80 and 443 respectively — when you installed IntelligenceCenter.

The following table describes each variable in the URL:

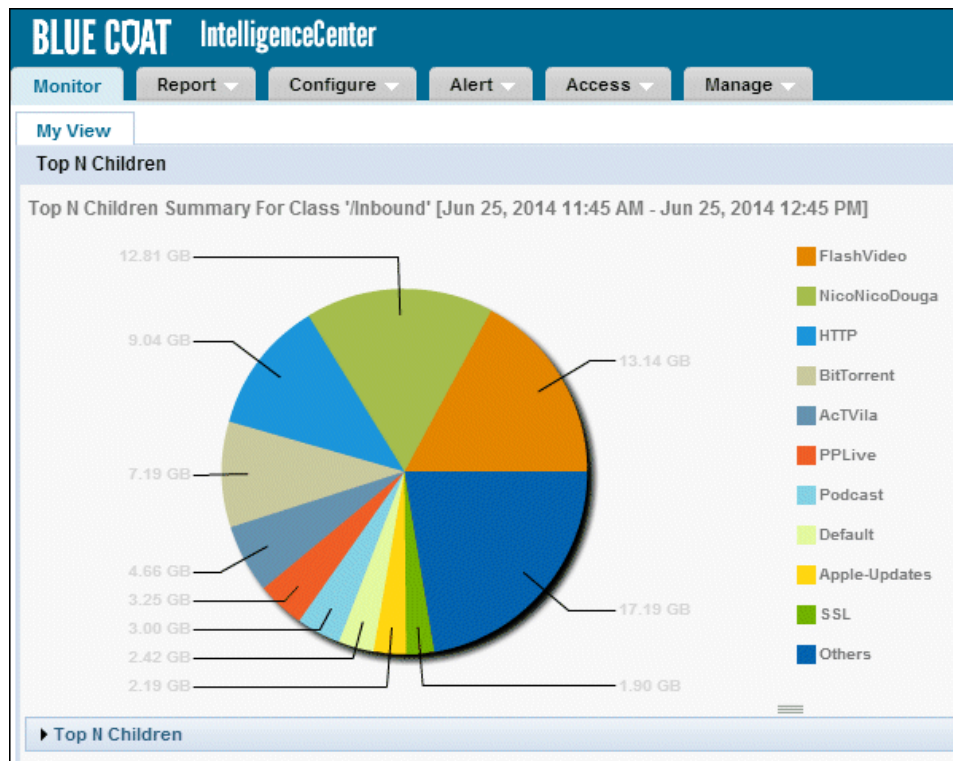
Variable	Description
<ipaddress>	<ul style="list-style-type: none">○ If you're logging in from a remote client, type the IP address or domain name of the computer on which you installed IntelligenceCenter. For example, <i>http://10.10.1.2</i>.○ If you're logging in from the local system, enter localhost. For example, <i>http://localhost</i>.

<port>	<ul style="list-style-type: none">○ If you accepted the default port numbers for HTTP (80) and HTTPS (443) when you installed IntelligenceCenter, you do not need to include the port number in the URL. For example, <i>http://10.10.1.2</i>.○ If you specified port numbers other than the defaults when you installed IntelligenceCenter, you must include them in the URL. For example: <i>https://10.10.10.1:8443</i> or <i>http://localhost:8080</i>.
--------	--

The login screen appears.

The image shows the login screen for Blue Coat IntelligenceCenter. It has a solid blue background. In the top left corner, the text "BLUE COAT" is in large, white, bold, sans-serif capital letters, followed by "IntelligenceCenter" in a smaller, white, sans-serif font. In the center of the screen, there is a white rectangular box containing the login form. Inside this box, the text "User Name" is to the left of a white input field. Below that, the text "Password" is to the left of another white input field. At the bottom of the input fields, there is a grey button with the word "Login" in white text.

3. Enter the **User Name** associated with your IntelligenceCenter [user profile](#). Your user profile configuration determines what IntelligenceCenter tasks you can perform.
4. Enter the **Password** that is associated with your user profile.
5. Click **Login**. The **Monitor** tab in IntelligenceCenter initially displays (but this is [configurable](#)).



Note: Your IntelligenceCenter session will automatically time out after 30 minutes of server inactivity. If your session times out, you will have to enter your password to restore your session.

Log Out of IntelligenceCenter

To log out of IntelligenceCenter:

Click the **Logout** link in the IntelligenceCenter banner to log out of the system.



The login screen is displayed. You can close the browser window or tab.

Note: Your IntelligenceCenter session will automatically time out after 30 minutes of server inactivity. If your session times out, you will have to enter your password to restore your session.

Manage Audit Logging

The audit log tracks user logins, logouts, configuration changes, portlet configurations, password changes, and other IntelligenceCenter activities such as the addition or removal of a device or user profile.

To view the audit log:

1. Select **Manage > Audit**. The audit log is displayed in the *Audit* pane.

Time	User	Description
▼ Today		
Dec 2, 2008 5:16 PM	admin	The User 'admin' logged in
Dec 2, 2008 4:29 PM	admin	The User 'admin' logged out
Dec 2, 2008 4:28 PM	admin	The User 'admin' logged in
Dec 2, 2008 4:28 PM	admin	The User 'admin' logged out
Dec 2, 2008 4:23 PM	admin	The User 'admin' logged in
Dec 2, 2008 11:35 AM	admin	The User 'admin' logged in
Dec 2, 2008 11:21 AM	admin	The User 'admin' logged out
Dec 2, 2008 11:08 AM	admin	The User 'admin' logged in
Dec 2, 2008 9:36 AM	admin	The User 'admin' logged in
Dec 2, 2008 1:31 AM	admin	The User 'admin' logged out
Dec 2, 2008 12:50 AM	system	The report definition 'Top N Traffic Classes Summary' was executed
Dec 2, 2008 12:40 AM	system	The report definition 'Top N Applications' was executed
Dec 2, 2008 12:30 AM	system	The DataCollector with the serial number '950-39715147' has been configured to start
Dec 2, 2008 12:27 AM	admin	The Person 'Timothy Andrew' was created
Dec 2, 2008 12:20 AM	system	The report definition 'Top N Listeners' was executed
Dec 2, 2008 12:10 AM	system	The report definition 'Top N Talkers' was executed

By default, all audit log entries that were generated today are displayed. For each entry, the log displays:

- The date and time of the log entry
 - IntelligenceCenter user who performed the activity
 - A description of the activity (such as *The User 'john' logged in* and *The PortletInstance 'Application Performance' was created*)
2. If you want to adjust the column widths, position the mouse pointer on the vertical line between column headings and drag to the right or left.
 3. If you want to change how the log entries are grouped, select a value from the **Arrange By** drop-down list. You can group log entries by **Time** (the default), **User**, or **Description**.
 4. If you want to filter the list of log entries that is displayed or to search for a specific entry:
 - Select the range of log entries you want to display by selecting a value from the **From** field drop-down list. You can choose to display log entries from **Today** (the default), **This Week**, **Last Week**, **This Month**, **Last Month**, **This Quarter**, or **Last Quarter** or you can choose **Other** and select a default time range. If you want to display log entries matching a specific search term, enter the term in the text box. Note that IC searches the **User** and **Description** fields only.
 - To search for log entries with a specific text string, enter the string in the text box.
 - Click **Find**. The audit log is updated to display only those entries that match the criteria you entered. To clear the filter, click **Clear**.
 5. If you want to sort the audit log, click the column heading by which to sort. An icon next to the heading indicates the sort column:

▲ indicates ascending order

▼ indicates descending order

Note: The sort order is restored to the default setting and the filter is cleared each time you return to the *Audit* pane. You can also clear the filter by clicking **Clear**.

Manage Licensing

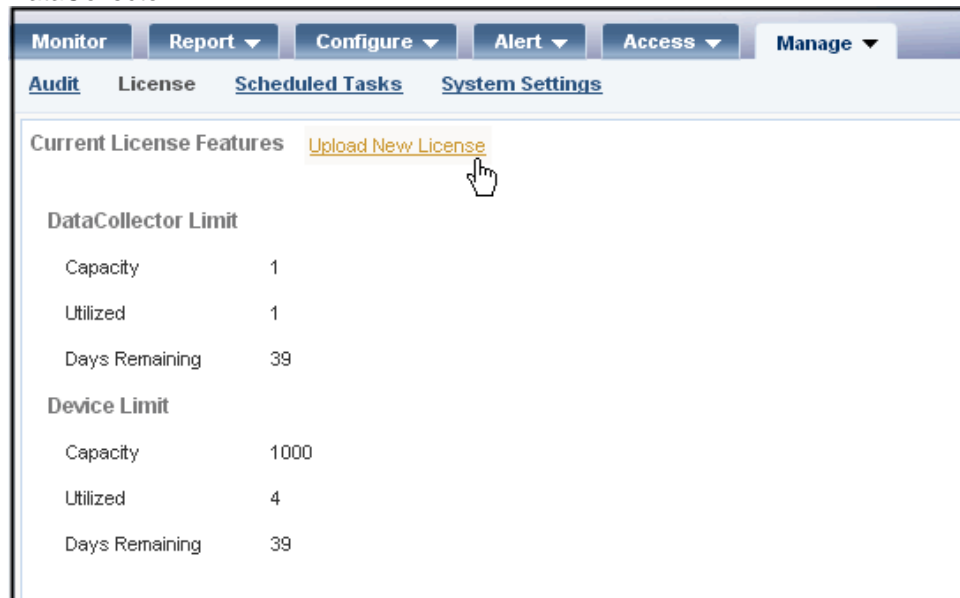
Your IC license controls which product features you are allowed to use. When you first install IC a 60-day trial license gets installed, which allows you to deploy a single DC and 1000 devices (the trial license does not support multiple DCs). To continue using IC after the 60-day trial period has ended, you must purchase the license components you require and then upload the corresponding license keys to IC. At a minimum you will need a base license, which allows you to deploy IC with a single DataCollector, as well as license keys to support the number of devices you plan to deploy. You must also purchase a license for each additional DC you plan to deploy.

After your purchase is complete, you will receive an email containing the activation codes for the licensing features you purchased. You must then go to Blue Touch Online (BTO) and download your license keys. When you download your license key file from Blue Coat, make sure you copy it to a location that is accessible by the system on which you installed IC. You can then use IC to install the license key file as follows:

Note: Do NOT open the license key in a text editor or rename it as this will corrupt the file. The license key must have a .bluecoat extension.

1. Go to <https://support.bluecoat.com/licensing> and select **License Others**. The Blue Coat Licensing Portal (BCLP) login screen is displayed.
2. Enter your Blue Touch Online (BTO) **User ID** and **Password**. If you do not have a BTO login, send an email to customercare@bluecoat.com. You will receive a response within one business day. Contact Customer Care for forgotten username and passwords or submit a request at <https://support.bluecoat.com/help/forgotpassword>. After you successfully log in, the BCLP Home page is displayed .
3. Enter the **Activation Code** contained in the email you received and then click **Next**. The download link for the corresponding license key is displayed. Download the file to a location that is accessible by the system on which you installed IC. Repeat this step for each activation code you received.
4. Select **Manage > License**. The *License* pane displays the following information about your device and DataCollector license:
 - **Capacity** — This section indicates the number of DataCollectors or devices your current license supports. In this release, a DataCollector **Capacity** value of 1 indicates that DataCollector is enabled; a value of 0 indicates that you are not licensed to use DataCollector. By default you are licensed to use a single DataCollector for 60 days.
 - **Utilized** — Indicates the number of devices or DataCollectors that you have currently deployed.
 - **Days Remaining** — Indicates the date and time at which your current license expires. The default DataCollector license expires 60 days after you deploy

DataCollector.



5. To install a new license, click **Upload New License**. A file selector dialog box is displayed.
6. Browse to the location of the license file and click **Open**. The new license features are reflected on the *License* window.
7. Click **Save** to accept the new license features.


Manage Scheduled Tasks


In IntelligenceCenter, scheduled tasks represent the static reports that you have scheduled (or reports that are scheduled run by default). You can view the scheduled reports or remove a scheduled report if it is no longer needed or if you want to change the configuration.

To view and/or delete reports that were previously scheduled:

1. Select **Manage > Scheduled Tasks**. A list of all the scheduled reports is displayed (including [custom reports](#)). Note that this list includes reports that you have scheduled as well as reports that were scheduled to run by default.


Scheduled Tasks				
	Top N Applications_System Weekly <i>Report process for the weekly TopNApplications</i>	Waiting	Last Run: Nov 16, 2008 12:35 AM	Next Run: Nov 23, 2008 12:35 AM
	Top N Applications_System Monthly <i>Report process for the monthly TopNApplications</i>	Waiting		Next Run: Dec 1, 2008 1:30 AM
	Top N Talkers_System Yearly <i>Report process for the yearly TopNTalkers</i>	Waiting		Next Run: Jan 1, 2009 2:00 AM
	Top N Listeners_System Weekly <i>Report process for the weekly TopNListeners</i>	Waiting	Last Run: Nov 16, 2008 12:15 AM	Next Run: Nov 23, 2008 12:15 AM
	Top N Applications_System Daily <i>Report process for the daily TopNApplications</i>	Waiting	Last Run: Nov 20, 2008 12:40 AM	Next Run: Nov 21, 2008 12:40 AM
	Top N Traffic Classes Summary_System Daily <i>Report process for the daily TopNTrafficClasses</i>	Waiting	Last Run: Nov 20, 2008 12:50 AM	Next Run: Nov 21, 2008 12:50 AM
	Top N Traffic Classes Summary_System Yearly <i>Report process for the yearly TopNTrafficClasses</i>	Waiting		Next Run: Jan 1, 2009 2:40 AM

2. To view details about a specific scheduled report, click the right-arrow  icon next to the report entry. The report entry expands to show detailed information about the report and the schedule on which it is configured to run.

 Top N Applications_System Weekly <i>Report process for the weekly TopNApplications</i>		Waiting	Last Run: Nov 16, 2008 12:35 AM	Next Run: Nov 23, 2008 12:35 AM
Schedule Details		Parameter Values		
Start	Nov 9, 2008 12:00 AM	Max Count	10	
End	No end date	Unit of Measure	bytes	
Recurrence:	Every week on Sunday	Sort Column	bytes	
		Sort Direction	desc	
		Show Guaranteed Rate	0	
		Failures		
Deactivate Delete				

3. If you want to stop a scheduled report from running, do one of the following:
 - To permanently remove the scheduled report, click the **Delete** link in the expanded report entry section.
 - To temporarily stop a scheduled report from running, click the **Deactivate** link. When you want to resume running the scheduled report, you can go back in and click the **Activate** link.

Note: If you want to change the configuration of a scheduled report, you must delete it and then [schedule the report](#) again.

4. To hide the schedule details, click the down-arrow  icon next to the report entry.

Configure System Settings

The following topics describe how to configure settings that control the operation of IntelligenceCenter:

- [Set device synchronization intervals](#)
- [Configure external syslog servers](#)
- [Define a login message](#)
- [Enable SNMP](#)
- [Define the system email address](#)
- [Enable external authentication](#)

Set Device Communication Intervals

The device communication intervals define how often IC polls devices for updated status and property information. By default, IC updates device status information every five minutes and device properties daily. Keep in mind that the shorter the intervals, the more network traffic and CPU usage increase.

To set the device communication intervals:

1. Select **Manage > System Settings > Device Communication**. The Device Communication settings are displayed.

Device Communication

Device communication settings define how often IntelligenceCenter gets updates from devices.

Update Device Status Every * 15 Minutes *Minimum 5 minutes*

Synchronize Device Properties




☐ Hourly

☐ Daily

☒ Weekly

☐ Monthly

Save Cancel

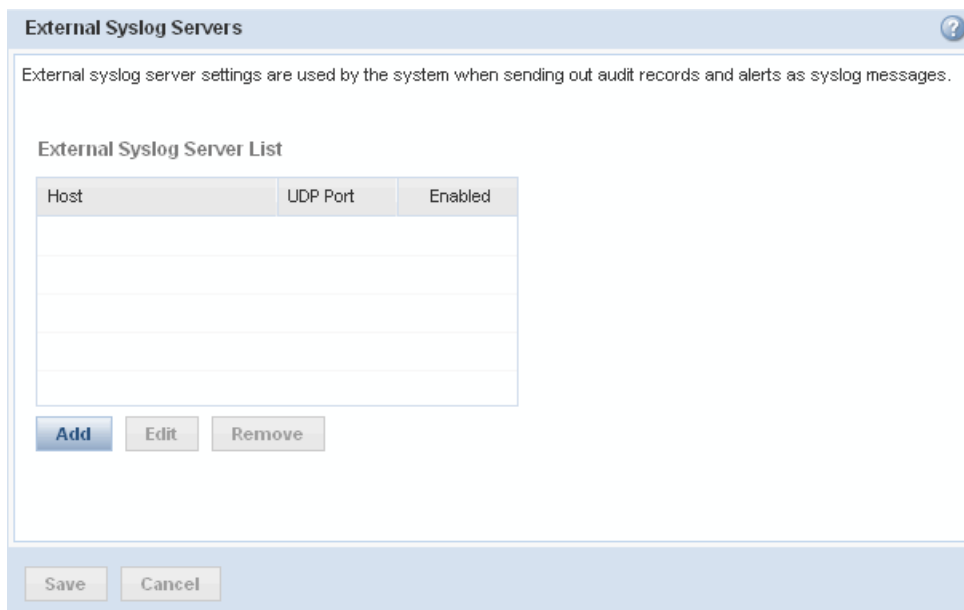
2. To modify the interval at which IC polls devices for connectivity, enter a new value (in minutes) in the **Update Device Status Every** text box. The resulting device status is indicated by a colored icon next to the device IP address or hostname on the **Devices** tab: a green  icon means that IC can ping the device and log in to it; a yellow  icon means that IC can ping the device, but cannot log in to it; a red  icon indicates that IC can neither ping nor log in to the device.
3. To modify the frequency at which IC should **Synchronize Device Properties**, select a radio button: **Hourly**, **Daily** (the default), **Weekly**, or **Monthly**. The [device property information](#) is displayed when you select a device on the **Devices** tab.
4. When you are done modifying the device communication settings, click **Save**.

Configure External Syslog Servers

IC can send audit log records and alerts generated by IC or DC to one or more external syslog servers. To enable alert notifications to be sent to a syslog server, you must configure IC to communicate with an external syslog server. In addition, you must [enable syslog within the alerting system](#).

To configure communication with external syslog servers:

1. Select **Manage > System Settings > External Syslog Server**. The External Syslog Servers pane is displayed.



External Syslog Servers

External syslog server settings are used by the system when sending out audit records and alerts as syslog messages.

External Syslog Server List

Host	UDP Port	Enabled

Add **Edit** **Remove**

Save **Cancel**

2. Click **Add**. The *Add External Syslog Server* dialog box is displayed.
3. Enter the IP address or DNS hostname of the external syslog server in the **Host** field.
4. If necessary, change the **UDP Port** (514 is the default port number for syslog).
5. To enable IC to send records to this syslog server, make sure the **Enabled** checkbox is selected.
6. Click **Save**. The external syslog server you added is displayed on the **External Syslog Server List**.
7. If you want to modify or remove external syslog server that you previously added, do the following:

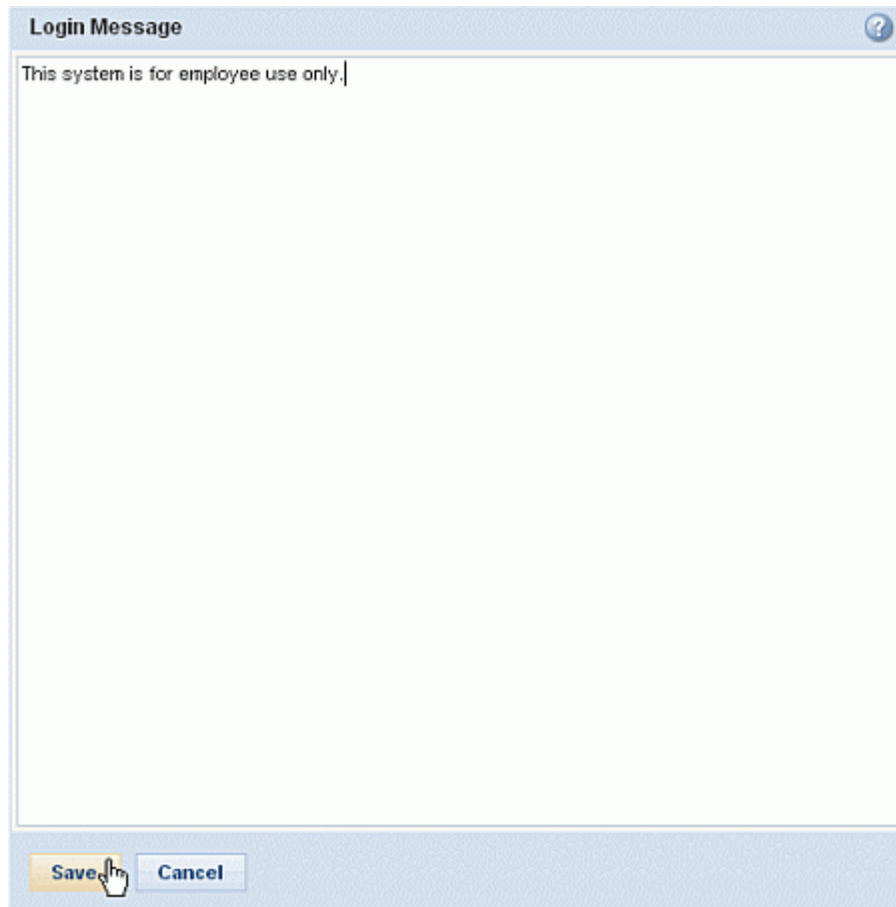
- To enable or disable communication with an existing syslog server, check or uncheck the corresponding **Enabled** checkbox in the **External Syslog Server List** and then click **Save**.
- To edit the configuration of an existing external syslog server, select the corresponding entry in the **External Syslog Server List** and click **Edit**. Follow Steps 3 through 6 to modify and save the configuration settings.
- To permanently remove an external syslog server that you previously added, select the corresponding entry in the **External Syslog Server List** and click **Remove**. Note that if you only want to disable communication with this server temporarily, you can simply uncheck the **Enabled** checkbox; then you can recheck it when you want to resume communication with this server.

Define the Login Message

The login message appears on the IntelligenceCenter login screen. This feature is useful for informing users about the company's access policies and consequences for unauthorized use.

To define a login message:

1. Select **Manage > System Settings > Login Message**. The *Login Message* pane is displayed.
2. Type the login message you want to display — up to 2048 characters long — in the **Login Message** text box.

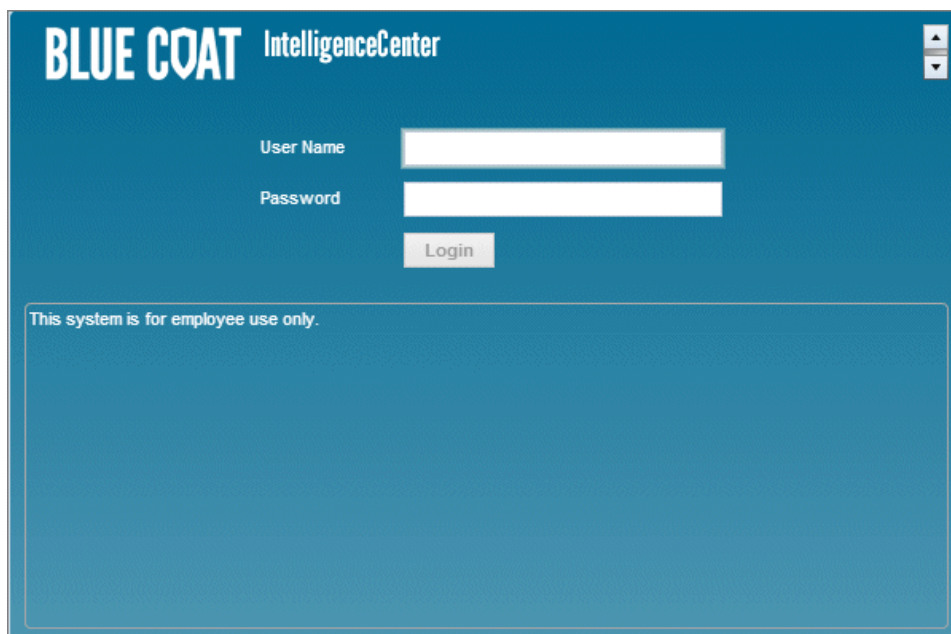
A screenshot of a web-based configuration window titled "Login Message". The window has a light blue header bar with the title and a help icon. Below the header is a large text area with a light gray background. The text "This system is for employee use only." is entered in the text area, with a cursor at the end. At the bottom of the window, there are two buttons: "Save" (highlighted with a mouse cursor) and "Cancel".

3. Click **Save**.

4. To see how the message looks on the login screen, click the **Logout** link in the banner.



The login screen will then display with your login message.

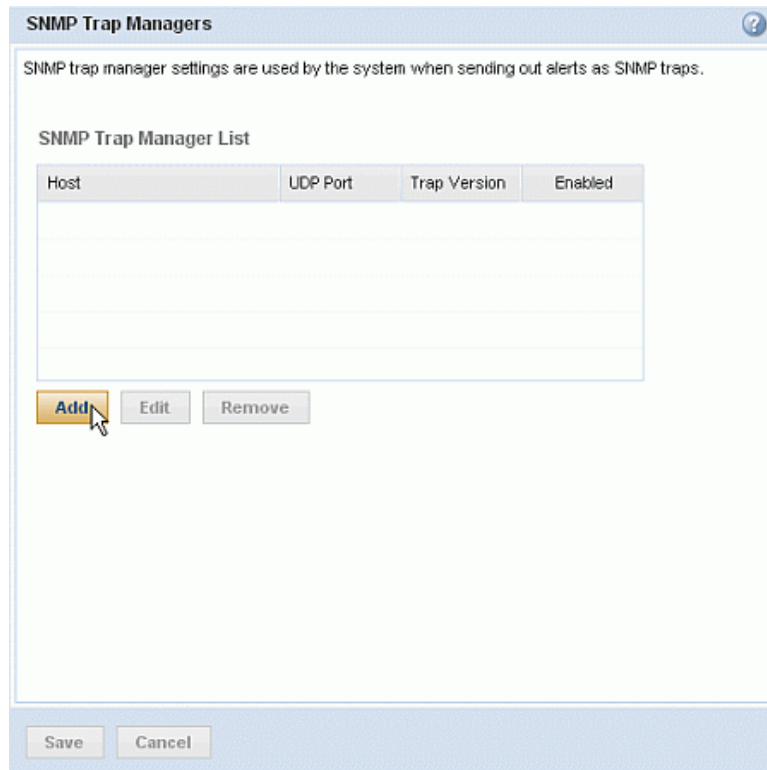
The login screen for Blue Coat IntelligenceCenter. It has a dark blue header with the "BLUE COAT IntelligenceCenter" logo. Below the header, there are two input fields for "User Name" and "Password", and a "Login" button. A large, empty rectangular box with a light blue border is at the bottom, containing the text "This system is for employee use only." in the top left corner.

Enable SNMP

IC can send alerts generated by IC or DC to one or more external SNMP trap managers. To enable this feature, you must configure the SNMP agent on IC to send SNMP traps to an external SNMP manager. In addition, you must [enable SNMP traps within the alerting system](#).

To configure communication with an SNMP trap manager:

1. Select **Manage > System Settings > SNMP Trap Managers**. The SNMP Trap Managers pane is displayed.



The image shows a software window titled "SNMP Trap Managers" with a help icon in the top right corner. Below the title bar, a text box states: "SNMP trap manager settings are used by the system when sending out alerts as SNMP traps." Below this is a section titled "SNMP Trap Manager List" containing a table with four columns: "Host", "UDP Port", "Trap Version", and "Enabled". The table is currently empty. Below the table are three buttons: "Add" (highlighted with a mouse cursor), "Edit", and "Remove". At the bottom of the window are "Save" and "Cancel" buttons.

Host	UDP Port	Trap Version	Enabled

2. Click **Add**. The *Add SNMP Trap Manager* dialog box is displayed.

3. Enter the IP address or DNS hostname of the SNMP trap manager in the **Host** field.
4. If necessary, change the **UDP Port** (162 is the default port number for SNMP traps).
5. Select the **SNMP Trap Version**, which dictates the SNMP security model that is used in exchanges between the IC SNMP agent and the SNMP manager. After you make a selection, the corresponding trap configuration fields are displayed at the bottom of the dialog box:
 - SNMPv1 or SNMPv2c do not provide authentication or encryption; all that is required is the **Community String**, which is like a password. The default community string, *public*, is used by many SNMP managers. However, you must use the community string expected by the specific SNMP manager you are configuring.
 - SNMPv3 requires a username. In addition, you can optionally choose a security model that requires authentication (authNoPriv) or a combination of authentication and encryption (authPriv) for added message integrity and privacy with SNMPv3. [Click here for a description of the fields for an SNMPv3 configuration.](#)

Field	Description
User name	The user name for the SNMP agent running on the IC server.
Authentication Mode	Select the authentication method to use when communicating with this trap manager: None , MD5 (message digest), or SHA (Secure Hash Algorithm).
Authentication Passphrase	Specify the authentication password; retype it in the

	Retype Passphrase field.
Privacy Mode	Select the encryption mode to use when communicating with this trap manager: None , AES (Advanced Encryption Standard), or DES (Data Encryption Standard).
Privacy Passphrase	Specify the encryption password; retype it in the Retype Passphrase field.

6. To enable IC to send traps to this SNMP manager, make sure the **Enabled** checkbox is selected.
7. Click **Save**. The SNMP trap manager you added is displayed on the **SNMP Trap Manager List**.
8. If you want to modify or remove an SNMP trap manager that you previously added, do the following:
 - To enable or disable communication with an existing SNMP trap manager, check or uncheck the corresponding **Enabled** checkbox in the **SNMP Trap Manager List** and then click **Save**.
 - To edit the configuration of an existing SNMP trap manager, select the corresponding entry in the **SNMP Trap Manager List** and click **Edit**. Follow Steps 3 through 7 to modify and save the configuration settings.
 - To permanently remove an SNMP trap manager that you previously added, select the corresponding entry in the **SNMP Trap Manager List** and click **Remove**. Note that if you only want to disable communication with this SNMP manager temporarily, you can simply uncheck the **Enabled** checkbox; then you can recheck it when you want to resume communication with this manager.

Define the System Email Address

To enable IC to send email notification of events—such as report generation notification—to selected users, you must configure an email server. You will need the DNS name or IP address of the Simple Mail Transfer Protocol (SMTP) server through which IC will send the email as well as the email address and an account for IC to use to send email from.

To set up the IC email account:

1. Select **Manage > System Settings > System Email Address**. The IC system email account settings are displayed in the right pane. You must replace the default values with actual email server and addressing information before IC will be able to send email notifications.

The screenshot shows a web interface with a left sidebar titled "System Settings" containing a list of configuration categories: Device Communication, External Syslog Servers, Login Message, SNMP Trap Managers, System Email Address (highlighted), and External Authentication. The main content area is titled "System Email Address" and contains a descriptive sentence: "System email settings define the settings IntelligenceCenter uses to email reports and alerts." Below this, there are three labeled input fields: "Email Server" with a red asterisk icon and the value "mail.acme.com", "Sender Email" with a red asterisk icon and the value "admin@acme.com", and "Sender Name" with a red asterisk icon and the value "IntelligenceCenter". At the bottom of the main area are two buttons: "Save" (highlighted with a mouse cursor) and "Cancel".

2. In the **Email Server** field, enter the fully qualified DNS name or IP address of your SMTP server.
3. In the **Sender Email** field, enter the email address of the user from which you want IC to send email notifications. This must be a valid email account on the specified email server.
4. In the **Sender Name** field, enter the name of the user from which you want IC to send email notifications.
5. Click **Save**.

Manage Heartbeats

IntelligenceCenter sends anonymous information (called *heartbeats*) to Blue Coat to assist Blue Coat personnel in analyzing how IntelligenceCenter is used. This basic information includes your license, server specifications, system resource use, and IntelligenceCenter configuration and use. No private, enterprise-sensitive data is transmitted. Although this feature is enabled by default, you have the option to decline sending heartbeats.

To enable or disable heartbeats:

1. Select **Manage > Heartbeat**. Information about the IntelligenceCenter Improvement Program is displayed.

IntelligenceCenter periodically checks for new versions.
Last Update Check: Aug 2, 2011 12:30 PM [Check Now](#)

IntelligenceCenter Improvement Program

This program submits anonymous usage statistics to Blue Coat using a secure (HTTPS) connection. The data assists Blue Coat with its ongoing IntelligenceCenter application reliability, quality and performance efforts.

Blue Coat collects:

- System information such as OS, platform, disk storage, RAM.
- IntelligenceCenter application information such as : IntelligenceCenter version, DataCollector versions.
- IntelligenceCenter application usage such as : number of DataCollectors, number of PacketShaper devices, LDAP etc.
- IntelligenceCenter application feature usage such as: Reports and Portlets usage, Business Hours Setup usage etc.

Blue Coat does not collect personal, enterprise-sensitive information such as hosts, traffic classes, network traffic data, users, passwords, groups or e-mails.

☒ Yes, my company will (anonymously) participate in the IntelligenceCenter Improvement Program.

☐ No, at this time my company declines to participate in the IntelligenceCenter Improvement Program.

[Example System Data](#)

[Save](#) [Reset](#)

2. Specify whether to send heartbeat information to Blue Coat. By default, heartbeats are enabled:
 - To disable heartbeats, select No, at this time my company declines to participate in the IntelligenceCenter Improvement Program.
 - To enable heartbeats, select Yes, my company will (anonymously) participate in the IntelligenceCenter Improvement Program.

Note: to see a sample heartbeat message, click **Example System Data**.

3. Click **Save**.

Manage Services

IntelligenceCenter and DataCollector run as Windows services on the systems on which they are installed. As part of your administration of IntelligenceCenter, you should monitor the services and, when necessary, restart them. Additionally, the IntelligenceCenter and DataCollector services will not run if the PostgreSQL services are not running.

To manage the services:

1. Click the **Start** button and select **All Programs > Administrative Tools**.
2. Choose **Services**.
3. Locate the following services:
 - PostgreSQL Database Server 8.4 (on the system where DataCollector and IntelligenceCenter are installed)
 - Blue Coat Agent For DataCollector 3.3.1.1 (on the system where DataCollector is installed)
 - Blue Coat DataCollector 3.3.1.1 (on the system where DataCollector is installed)
 - Blue Coat IntelligenceCenter 3.3.1.1 (on the system where IntelligenceCenter is installed)
4. Make sure the *Status* of the services is *Started* and the *Startup Type* is *Automatic*. To change these settings, you can right-click.
5. To restart a service:
 - Double-click the service you want to restart — PostgreSQL Database Server 8.4, Blue Coat Agent For DataCollector 3.3.1.1, Blue Coat DataCollector 3.3.1.1, or Blue Coat IntelligenceCenter 3.3.1.1 — and then click **Stop**.
 - After the service successfully stops, click **Start** and then click **OK** to close the dialog box.

Troubleshoot IntelligenceCenter

If you are experiencing difficulties while using IntelligenceCenter, follow these steps:

Step 1:

[Log out](#) of IntelligenceCenter (or just close your browser window or tab) and then try [logging in](#) again. If this doesn't solve your problem, continue to Step 2.

Step 2:

Ensure that the IntelligenceCenter and/or DataCollector [services](#) are running properly. If restarting the services doesn't solve your problem, continue to Step 3.

Step 3:

Reboot the system on which IntelligenceCenter and/or DataCollector are installed.

Step 4:

Check the IntelligenceCenter log files:

Log File	Location
IntelligenceCenter Application Log	<i><IntelligenceCenter Install Location>/logs/ic.log</i>
DataCollector Agent Log	<i><DataCollector Install Location>/logs/agent/collector-agent.log</i>
DataCollector Log	<i><DataCollector Install Location>/log/collector_<Day>.log</i>

Manage the Databases

IC and DC each have their own PostgreSQL databases that get created when you initially install the product. On IC, the database—called *PolicyVisionDB*—is used to store configuration settings, archived reports, audit logs, user profile settings, device information, report schedules, discovered traffic classes, and licensing information. On DC, the database—called *collector*—is used to store all of the metric (ME) and flow detail record (FDR) data that it collects from the devices on your network. Depending on the number of devices you are collecting from and the type and amount of data you are collecting (flow data consumes much more disk space than ME data), the DC database can grow quite large. When the database reaches 95% of its disk space capacity, DC will stop collecting data and reporting will be interrupted. You can [monitor the amount of available disk space](#) on the drive that contains the DC database from the DataCollector **Status** tab so that you can take action to [free up disk space](#) on the drive before collection stops.

Because these databases contain all of your configuration settings and all of your data, it is important that you perform regular backups of both databases. This way, if you experience a system failure, you will be able to restore your IC and DC databases, either on the same system or on a new system, without losing all of your configuration settings and historical data. As a security measure, you should also back up the databases before you perform a software upgrade.

Keep in mind that the IC and DC databases must stay in sync. Therefore, in order to perform a successful database restore—of either the IC or DC database—you must restore both databases. Additionally, to ensure that the databases stay in sync, you must restore both databases using backup files that were created at roughly the same time.

The following topics describe the backup and restore procedures:

- [Back Up the Databases](#)
- [Restore the Databases](#)

Back Up the Databases

You back up the IntelligenceCenter (IC) and DataCollector (DC) databases using the `pg_backup.py` script file. This script gets installed to the `C:\Program Files\PostgreSQL\8.4\tools` directory when you install IntelligenceCenter version 3.x.

If IC and DC are running on the same server, this script will automatically back up both databases. If IC and DC are running on separate servers, you must run the script on both systems. You **MUST** back up both components in order to be able to successfully restore them later.

To back up the database(s):

1. Open a Windows command shell.
2. Create the directory where you want to store the backup files. As an extra security measure, you may want to create the backup directory on a different system and/or disk drive than your IC and DC databases are installed on in case there is a system failure.
3. Change to the `C:\Program Files\PostgreSQL\8.4\tools` directory.
4. At the command prompt enter the following command:

```
python pg_backup.py -b <backup_location>
```

where `<backup_location>` is the path and name of the folder you created to store the backup files. For example:

```
python pg_backup.py -b H:\Database_backups
```

5. When the script finishes, the command prompt returns. You will find the following backup files in backup directory you specified when you ran the script:
 - `db_collector_yyyymmdd.dmp`—This is the backup for the DataCollector database. The `yyymmdd` represents the date that the backup was created. For example, if you ran the script on December 1, 2008, the file name would be `db_collector_20081201.dmp`. This file will only be present if DC is installed on the system.
 - `db_PolicyVisionDB_yyyymmdd.dmp`—This is the backup for the IntelligenceCenter database. The `yyymmdd` represents the date that the backup was created. For example, if you ran the script on December 1, 2008, the file name would be `db_PolicyVisionDB_20081201.dmp`. This file will only be present if IC is installed on the system.
6. If IC and DC are not installed on the same server, repeat this procedure on the server where the other component is installed. To ensure that the databases stay in sync, the backup files for IC and DC must be created at roughly the same time. If the components are installed on separate systems, make sure you run your backups at the same time or one after the other.

Restore the Databases

You restore the IntelligenceCenter (IC) and DataCollector (DC) databases using the `pg_restore.py` script file. This script gets installed to the `C:\Program Files\PostgreSQL\8.4\tools` directory when you install IntelligenceCenter. In order to restore the databases, you must have previously [backed them up](#) using the `pg_backup.py` script. If IC and DC are running on the same server, the `pg_restore.py` script will automatically restore both databases. If IC and DC are running on separate servers, you must run the script on both systems. You **MUST** restore both components in order for the databases to remain in sync.

Before you begin the restore operation, you must know the path and directory where your backup files (`db_collector_yyyymmdd.dmp` for DC and `db_PolicyVisionDB_yyyymmdd.dmp` for IC) are stored. You can restore the databases on an existing sever or you can restore to a new server. If you will be restoring the databases on a new system, make sure that the backup files are in a location that is accessible by the new system.

When the restore script runs, it automatically restores the database using the backup (`.dmp`) file with the latest date as indicated by the `yyymmdd` portion of the file name. Therefore, if you specifically want to restore a backup file other than the most recent one, you should remove any more recent backup files from the directory. Keep in mind, however, that in order for the databases to stay in sync, the backup files must have been created at roughly the same time. Therefore, you should not attempt to restore backups of IC and DC that were created on different days.

Note: Before you can run the restore script, you must stop all IC and DC services, interrupting access to the IntelligenceCenter application.

Restore the databases—either on the same server or on a new server—as follows:

1. Install IC and DC using the same configuration settings you used when you previously installed them (for example, make sure you use the same DC touch password), but **DO NOT** restart the server until after you restore the databases.
2. Verify that all IC and DC services are stopped:
 - Blue Coat Agent For DataCollector 3.3.x.x (on systems where DataCollector is installed)
 - Blue Coat DataCollector 3.3.x.x (on systems where DataCollector is installed)

- Blue Coat IntelligenceCenter 3.3.x.x (on systems where IntelligenceCenter is installed)
3. On the server where you want to restore the database(s), open a Windows command shell.
 4. Change to the C:\Program Files\PostgreSQL\8.4\tools directory on the server.
 5. At the command prompt enter the following command:

```
python pg_restore.py -b <backup_location>
```

where <backup_location> is the path and name of the folder where the backup files are located. For example:

```
python pg_restore.py -b H:\Database_backups
```

When the script finishes, the command prompt returns.

6. Restart the IC and DC services.
7. If IC and DC are not installed on the same server, repeat this procedure on the server where the other component is installed.

Manage Disk Space on the Data Drive

When the drive that is hosting the DC database reaches 95% of its capacity, DC will automatically stop collecting data. You can [monitor the disk space usage on your DC](#) from the DataCollector **Status** tab. If you are nearing the 95% mark, you should take steps to ensure that data collection is not interrupted.

First, ensure that the disk space is being consumed by the DC database and not by another application. To check the amount of disk that is being used for the DC database, use Windows Explorer to navigate to the directory where your DC data directory is located (C:\ProgramFiles\PostgreSQL\8.4\data by default). Right-click the data folder to display the **Properties** dialog box. On the **General** tab, check the **Size** field to determine the actual size of the DC database. If you find that the disk space is not being consumed by DC, you must free up the disk space that is being used by other applications so that DC can continue to collect data. If you determine that the DC database is consuming the disk space, you must reduce the amount of data DC collects and stores as follows:

- Reduce the amount of data that is stored in the DC database by [modifying the data retention values](#).
- Reduce amount of FDR that your PacketShaper devices send to DC by setting up FDR filters. You can filter the FDR records by class, service, and/or subnet using the `setup flowrecords filters` CLI command on each PacketShaper. For more information, refer to [setup flowrecords filters Collectors](#) in PacketGuide.
- Move the data directory to a larger drive. To do this, you must [back up the IC and DC databases](#) and then [restore them](#) to a different location.

Set Up Data Collection and Reporting

Before you can start generating IntelligenceCenter (IC) reports, you must configure your IC network. At a minimum, an IC network contains a group of devices and a DataCollector (DC) that collects and reports on the data generated by these devices. Depending on the amount of data you need to collect and the way you want to partition your reporting, you may need to create multiple groups, each with its own DC and unique set of devices. Each DC collects data from and reports on the devices that reside in its own group.

To set up your IC network, you must complete the following tasks:

1. Define the [groups](#) and [sub-groups](#) that will contain the DCs and devices in your IC network. Each group you create represents a standalone reporting domain with a unique DC and set of devices.
2. [Add a DC](#) to each group.
3. [Add each device](#) that you want to be able to report on to the IC network.
4. [Configure DC to collect data from the devices](#).
5. (optional) [Define a set of business hours](#) to allow you generate reports that only include data for the times when the majority of network activity occurs.
6. [Define the applications that you are interested in reporting on and designate critical applications](#).
7. [Enable site reporting by associating IC site names with PacketShaper location-based traffic classes](#).

Data Collection Overview

The PacketShaper appliances on your network measure hundreds of characteristics about traffic as it passes, creating an extensive collection of [metric data \(ME\)](#). Additionally, if you use the [Flow Detail Record \(FDR\)](#) feature on your PacketShaper appliances, the appliances also gather per-flow statistics, including source and destination IP addresses, the size of the flow (in terms of packets and bytes), and when the flow was sent and emits these FDR records to a data collector.

You can configure DataCollector to collect ME and/or FDR data from any PacketShaper or PacketShaper ISP appliance running PacketWise version 7.3.1 or higher or any PacketShaper S500 appliance running 11.1.1.14 or higher. You can then use IntelligenceCenter to generate reports that aggregate the data that DataCollector collects from the appliances throughout your network. Some IC reports and portlets are based on Metric data, some are based on FDR data, and some use both ME and FDR data. Therefore, if you do not collect both types of data, you will not be able to view all portlets and reports.

DataCollector can also collect FDR data from other devices on your network, such as routers and switches, and report on the data.

Note: Before you can begin collecting FDR from a PacketShaper appliance or from a network device such as a router or a bridge, you must configure the device or appliance to emit FDR to Data Collector. For instructions on how to configure a PacketShaper to emit FDR, refer to [Define Flow Detail Record Collectors](#) in PacketGuide. For instructions on how to configure a network device to emit FDR to DataCollector, refer to the documentation for the device.

After DataCollector collects the metric and FDR data, it grooms the data and adds it to its database tables in a "ready to report" format. As DataCollector collects data over a period of time, it rolls the data up into time-based tables to enable reporting over a variety of time frames.

DataCollector maintains tables for raw, hourly, daily, monthly, and yearly data. To maximize your disk space usage, you can [set data retention policies](#) for each type of data separately. For example, for metric data, you can set separate data retention policies for Class, Partition, and Link data. Similarly, for FDR data you can set data retention policies for raw, hourly, daily, monthly, and yearly data. In addition, you can configure the DataCollector to [back up the FDR data](#) to a comma-separated values (CSV) file that you can export into an external application for aggregation and manipulation.

Additionally, DataCollector discovers traffic classes and partitions on the appliances from which it is configured to collect and pushes this information out to IntelligenceCenter. After traffic classes have been discovered, you can report on the classes or on the [applications](#) that are based on the traffic classes.

Measurement Variables Overview

The PacketShaper appliances on your network measure hundreds of characteristics about traffic as it passes, creating an extensive collection of measurement engine (ME) data (also called *metric* data). IntelligenceCenter allows you to aggregate the metric data that is collected by the appliances throughout your network so that you can compare, correlate, and summarize network behavior throughout your organization.

The measurement data that DataCollector collects depends on the appliance. DataCollector collects measurement data from the following appliances:

- [PacketShaper](#)
- [PacketShaper ISP](#)

Flow Detail Records Overview

Flow Detail Records (FDRs) provide per-flow statistics for traffic passing through the devices on your network. When the FDR feature is enabled on a device such as a PacketShaper, router, bridge, switch, or hub, the device becomes an *emitter*, periodically pushing FDRs to a remote system called a *collector*. The device will emit records that contain details of all flows that go through it to [DataCollector](#). You can then use IntelligenceCenter to [view reports](#) to summarize and analyze the data.

Note: In order to collect and report on FDR data from within IntelligenceCenter, you must configure your network devices (including your PacketShaper appliances as well as other devices such as routers or bridges, you must configure the device or appliance to emit FDR to Data Collector. For instructions on how to configure a PacketShaper to emit FDR, refer to [Define Flow Detail Record Collectors](#) in PacketGuide for PacketWise. For instructions on how to configure a network device to emit FDR to DataCollector, refer to the documentation for the device.

Generally, an FDR contains information about a TCP or non-TCP flow, such as source and destination IP addresses, the size of the flow (in terms of packets and bytes), and when the flow was sent. The specific fields of information vary according to the type of record format. DataCollector collects two types of FDR records: Packeteer-2 (PacketShaper only) and Netflow-5.

The NetFlow-5 record type identifies the flow's Layer 4 protocol (such as TCP, UDP, or ICMP) and IP ToS/Diffserv values. The Packeteer-2 format contains all the NetFlow fields as well as PacketShaper-specific data, for example: the traffic class into which the flow was classified, type of policy, number of retransmitted bytes, Response Time Measurement (RTM) data, packet exchange time, and VoIP statistics for RTCP VoIP streams.


Add DataCollector to the Network Group

After you install DC and IC, you must add DC to the network group to finish configuring it.


Note: The default license that is included with IC allows you to deploy a single DC for 60 days. If you want to continue using DC after the 60-day trial period ends, you must purchase a [license](#) to do so. To deploy additional DCs you must obtain a license.

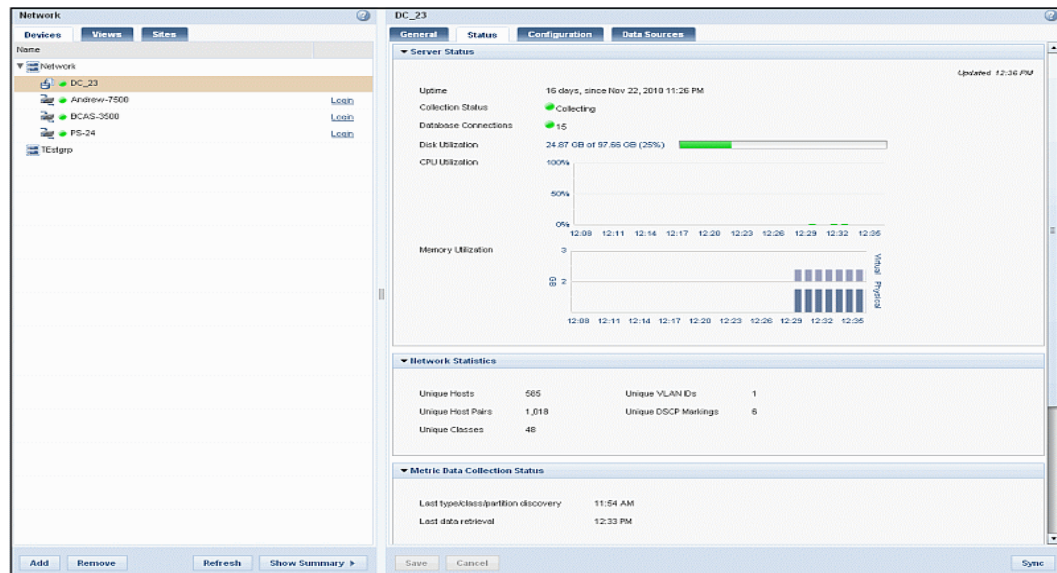
To add DataCollector to IntelligenceCenter:

1. Select **Configure > Network > Devices**.
2. Select the [network group](#) where you want to create DC. You must have one DC per network group.
3. Click **Add** and then select **DataCollector** from the pop-up menu.



4. The *Import DataCollector* dialog box appears.

5. Enter a **Device Name** for DC or accept the default name.
6. Enter an optional **Description** of the DataCollector.
7. In the **Host** field, enter the IP Address or hostname of the server on which DC is installed. If IC and DC are installed on the same system, use `localhost`. Note however that you can only run one DC on a given server; if you have multiple DCs, you must install them on separate servers, each of which meets the System Requirements as documented in the *IntelligenceCenter Getting Started Guide*.
8. Enter the DC **Touch Password**. This is the password that you created when you installed this DC.
9. In the **Location** field, you can optionally enter a description of the server's location.
10. Specify which group this DC will service. A DC can only collect and report on the devices contained in its associated group:
 - If the group that this DC will service already exists, select the **Select from current Group(s)** radio button and then select the group from the **Groups without DataCollectors** drop-down list.
 - If the group does not yet exist and you want to create it now, select the **Create new Group** radio button and then enter a **Name** and optionally a **Description** for the new group.
11. Click **Save** to save the DC configuration and the group, if applicable. A DC object appears on the **Devices** tab in the **Network** pane, as indicated by the  icon. The DC tabs appear in the right pane.




Define Data Sources

DataCollector is a software component that collects data from one or more network devices including PacketShaper appliances, routers, switches, gateways, hubs and/or bridges. The devices from which DataCollector collects are called data sources. You can collect data from any PacketShaper or PacketShaper ISP appliance running PacketWise version 7.3.1 or higher or any PacketShaper S500 appliance running 11.1.1.14 or higher. In addition you can collect flow detail record (FDR) data from NetFlow-capable devices such as routers and switches. Before you can define an appliance or device as a data source, you must [add it to the IC network](#).

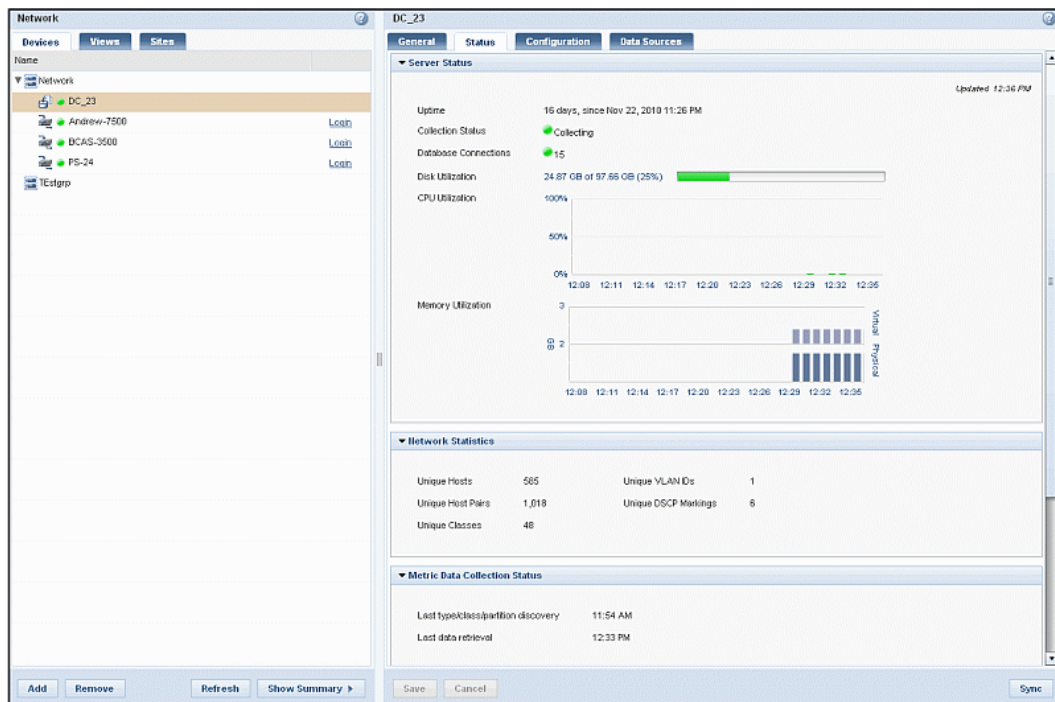
To enable DC to collect from a device, you must add it as a data source as described in the following procedure. When configuring a device as a data source, you also define the types of data to collect from it: [metric \(ME\) data](#) and/or [flow](#) data. Keep in mind that if plan to collect flow data from a data source, you must also configure the device to emit flow data to DataCollector. For instructions on how to configure a PacketShaper appliance to emit FDR, refer to [Define Flow Detail Record Collectors](#) in PacketGuide. For instructions on how to configure another type of network device to emit flow data, refer to the documentation for the device.

To add data sources to the DataCollector configuration:

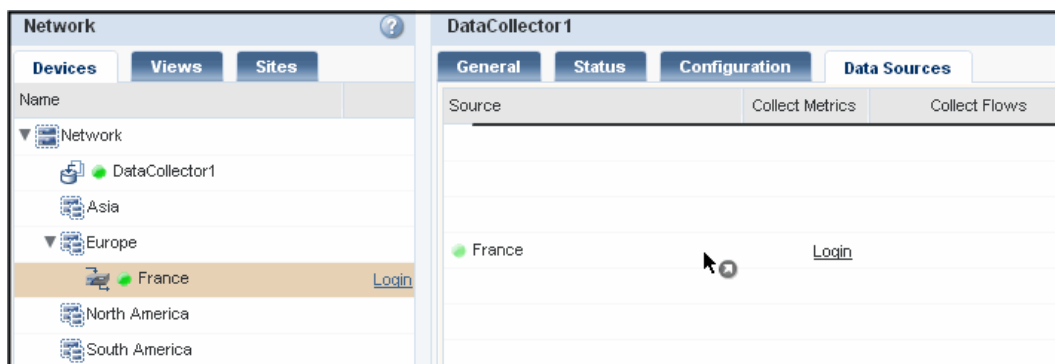
1. Select **Configure > Network > Devices**.
2. Select the DC object for which you want to add data sources. A DC is identified by the  icon.

Note: To expand or collapse a network group, click the arrow icons next to each group or sub-group. The down arrow ▼ icon indicates that the group is expanded; the right arrow ► icon indicates that the group is collapsed.

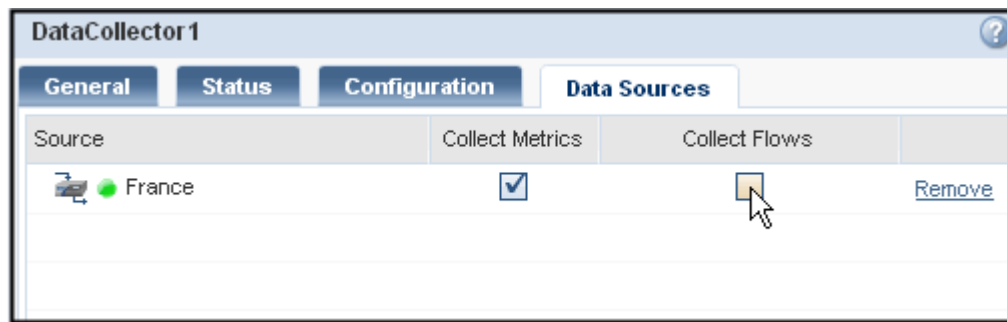
When you select the object, the DataCollector [Status tab](#) is displayed.



3. Select the **Data Sources** tab. Initially, there are no data sources configured.
4. In the **Devices** tab in the **Network** pane, locate the sub-group, view, or appliance that you want to add as a data source. IC allows you to define any level of the network hierarchy or any network view as a data source. Any PacketShaper appliances, and/or NetFlow-5 capable network devices (such as routers) in the network branch or view are then automatically added as data sources and use the data collection settings you define for the branch or view (ME, Packeteer-2, and/or NetFlow-5).
5. Select the group, sub-group, or appliance and drag it to the **Source** column in the **Data Sources** tab. Note that if you add a sub-group as a data source, any PacketShaper appliance that you add to the group or sub-group later will automatically be recognized as a data source. Keep in mind that this also means that if you remove an appliance from the group or sub-group— either moving it to another sub-group or removing it from the topology — DataCollector will no longer collect data from the appliance.



As soon as the branch, view, or appliance is successfully added as a data source, checkboxes appear in the **Collect Metrics** and **Collect Flows** columns of the corresponding row.



Note: Because you can add views or branches at any level of the hierarchy, an individual device may be added as a data source multiple times, each with different data collection settings (for example, if you add a parent branch and a child branch as separate data source entries; if you add a branch and an individual device within the branch; or if you add a branch and a view that contain overlapping devices). The settings you define at the child level will always take precedence over the settings defined at higher levels of the hierarchy. For example, suppose you want to collect metric data from all PacketShaper appliances in a branch and you want to collect flow data from a single PacketShaper in the branch. In this case, you could add the branch as a data source and configure it to collect metrics. You could then add the individual PacketShaper from which you want to collect FDR as a data source and configure it to collect flows.

6. Specify the types of data you want the DataCollector to collect from the new data source as follows:
 - If you want to collect [ME data](#) from the appliance, check the **Collect Metrics** checkbox. The specific ME variables that are collected depend on the appliance type.
 - If you want to collect [FDR data](#) from the appliance or network device, check the **Collect Flows** checkbox. Note that you must also configure the appliance to emit FDR to DataCollector. For instructions, refer to [Define Flow Detail Record Collectors](#) in PacketGuide. You can collect both [Packeteer-2](#) and [NetFlow-5](#) from a PacketShaper appliance; you can only collect NetFlow-5 from other types of network devices, such routers and switches.
7. If you want to add additional data sources, repeat steps 5 through 7.

8. When you are done adding data sources, click **Save**.

[illegible]

Set Up Business Hours Reporting

The business hours reporting feature allows you to define a set of hours that correspond to the hours when your network usage is most active (that is, the hours when people are working). You can define one set of business hours per DataCollector. After you define the business hours, you can then run your reports to show data for only those hours. This way, statistical information displayed in the reports is not skewed by the hours when there is very little traffic, such as weekends and evenings.

Note: If you have multiple DCs you can [share the business hours configuration](#) between multiple DCs using the import/export feature. This way you do not have to manually configure each DC.

To define business hours for a DataCollector:

1. Select **Configure > Network**.
2. On the **Devices** tab, select the DataCollector for which you want to specify business hours.
3. Select the **Business Hours Setup** tab.
4. Select the **Enable Business Hour Setup** checkbox.
5. Select the time zone that corresponds to the region where you are setting your business hours (this does not have to be the same as the time zone where the DataCollector resides) by selecting a value from the **Select Time Zone** drop-down list.

Business Hours settings provide flexibility to view reports based on the configured business hours.

Enable Business Hour Setup ☒

Select Time Zone (GMT+05:30) Chennai, Kolkata, Mumbai, New Delhi ▼

Work Time

General Work Hours

Start Time 09:00 ▼ Number of Working Hours 09:00 ▼

End Time 18:00 ▼

Work Week ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☐ Sat ☐ Sun

Import And Export Options

Import Export

Save Cancel Sync

6. Select the work hours as follows:
 - a. Select the **Start Time** from the drop-down list.

Note: Business hours can be set on the hour or the half hour only. In the case where the business hours time zone and business hours start time or end time boundary differ by 30 minute fraction, the reports will display data for the whole hour. For example, if you set business hours for GMT-8 from 8:30AM-5:30PM, the reports will show data for the time period from 8AM-6PM.

- b. Select the **Number for Working Hours** from the drop-down list. When you make a selection, the **End Time** is automatically calculated and displayed.
 - c. Select the days of the week that network users will be working by selecting the appropriate checkboxes in the Work Week field.
7. To save the business hours definition for this DataCollector, click **Save**.

Now that you have defined business hours, you can [run reports](#) that only include data for the specified hours.

Share Business Hours Settings Between Multiple DCs

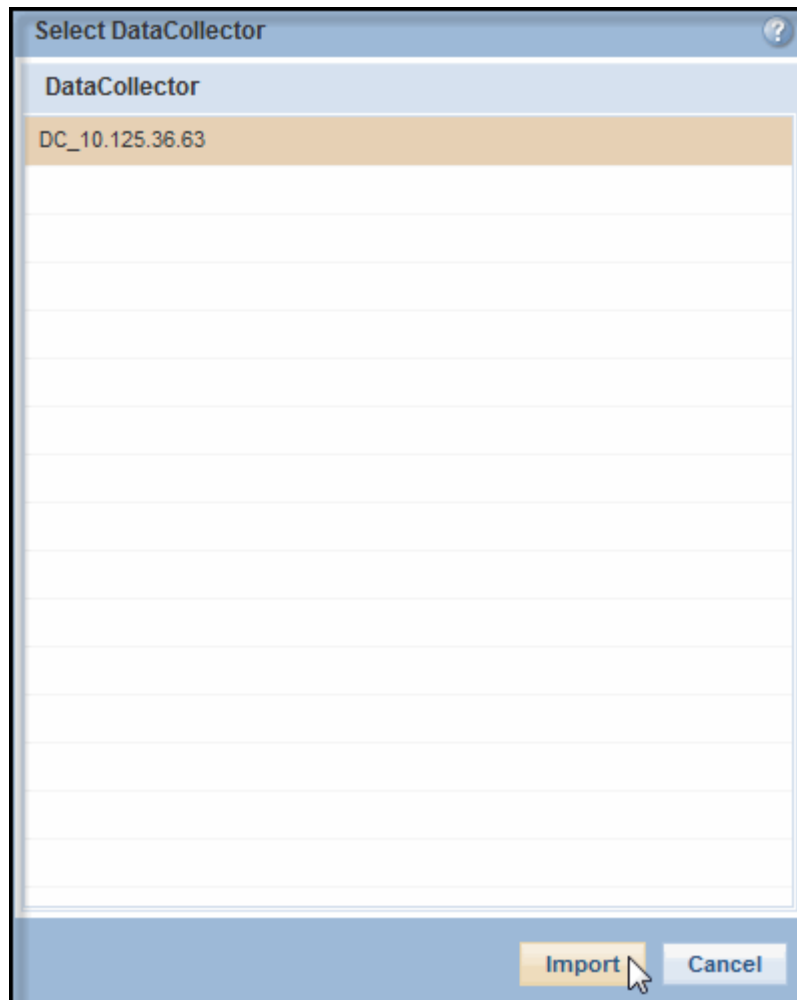
If you have multiple DCs you can share the business hours configuration between multiple DCs using the import/export feature. This way you do not have to manually configure each DC. You can:

- Import business hours settings from a DC you already configured onto an unconfigured DC.
- Export business hours settings from a configured DC to other, unconfigured DCs.

Importing Business Hours From Another DC

To import business hours settings from another DataCollector:

1. Select **Configure > Network**.
2. On the **Devices** tab, select the DataCollector on which you want to import the business hours settings.
3. Select the **Business Hours Setup** tab.
4. Select the **Enable Business Hour Setup** checkbox.
5. Click **Import**. The *Select DataCollector* dialog displays.
6. Select the DataCollector you want to import business hours from and then click **Import**.

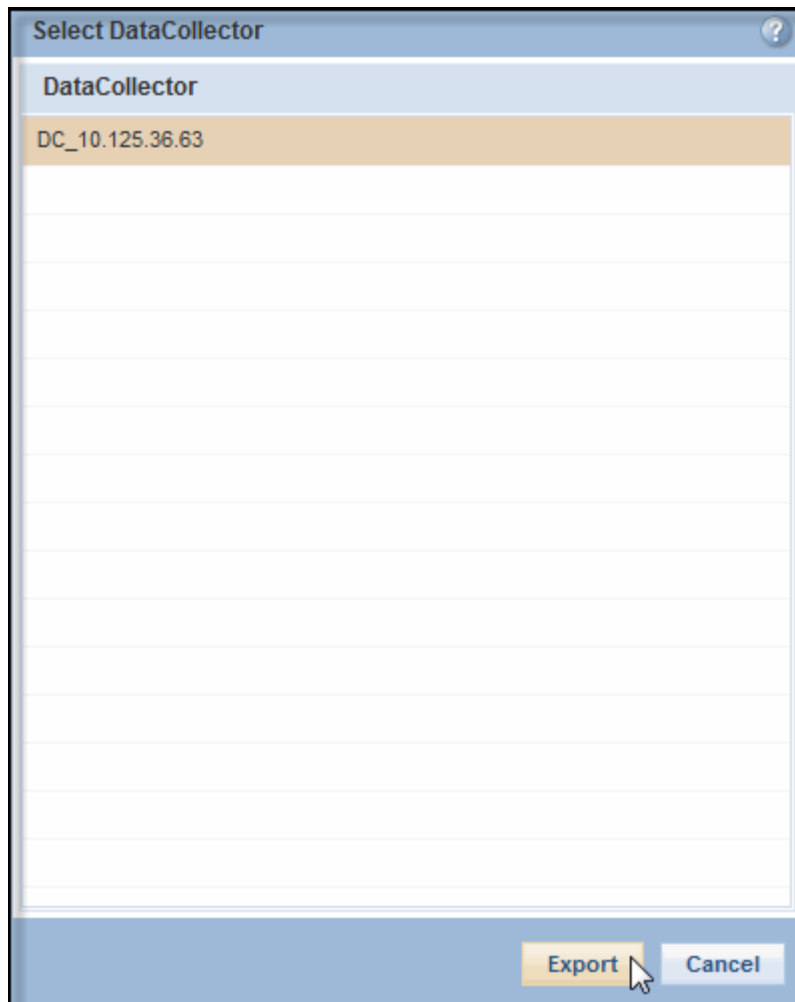


7. Click **Save**.

Exporting Business Hours to Another DC

To export the business hours settings from this DC to another DataCollector:

1. Select **Configure > Network**.
2. On the **Devices** tab, select the DataCollector on which you want to import the business hours settings.
3. Select the **Business Hours Setup** tab.
4. Select the **Enable Business Hour Setup** checkbox.
5. Click **Export**. The *Select DataCollector* dialog displays.
6. Select the DataCollector you want to import business hours from and then click **Export**.




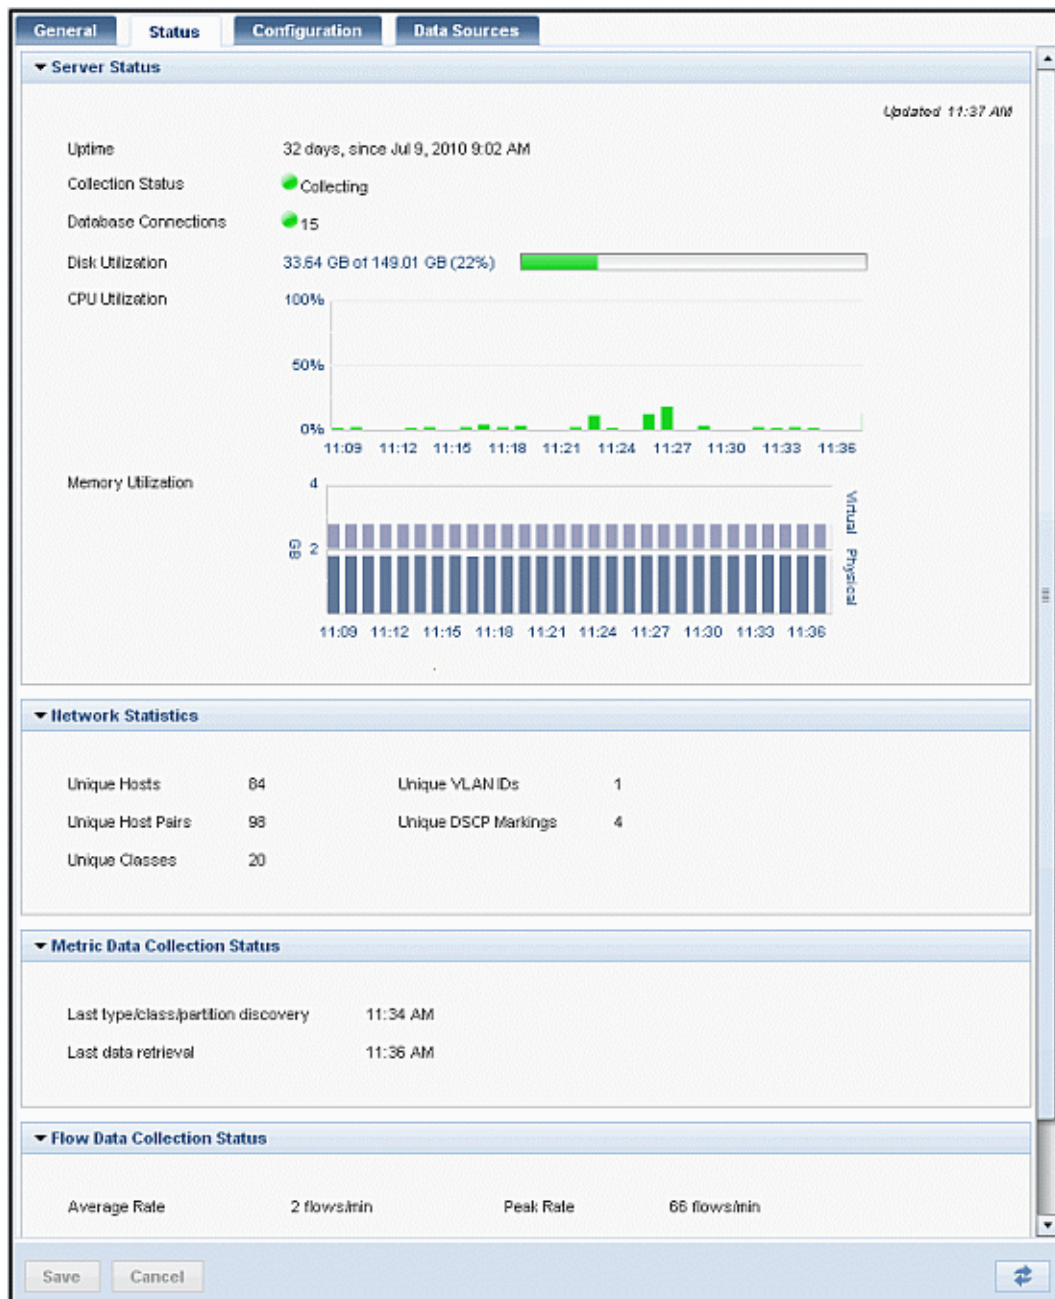
7. Click **Save**.

Force Class Discovery

By default, DataCollector sends an updated list of traffic classes to IntelligenceCenter every 24 hours. If you want to ensure that the list of traffic classes is up to date, you can force a class discovery.

To force a class discovery:

1. Select **Configure > Network > Devices**.
2. Select the DataCollector device entry (indicated by the  icon) in the network topology. When you select the object, the DataCollector **Status** tab is displayed in the right-hand pane.



3. Select the **General** tab.
4. Click **Discover Traffic Classes**. DataCollector pushes the current traffic class information it has collected from its data sources to IntelligenceCenter.

Note: If traffic class discovery occurred within the last 15 minutes, either automatically or manually, IC will not allow you to initiate another class discovery. In this case, wait for 15 minutes to ensure that the previous class discovery is complete and then try again.


Modify Data Collection Settings

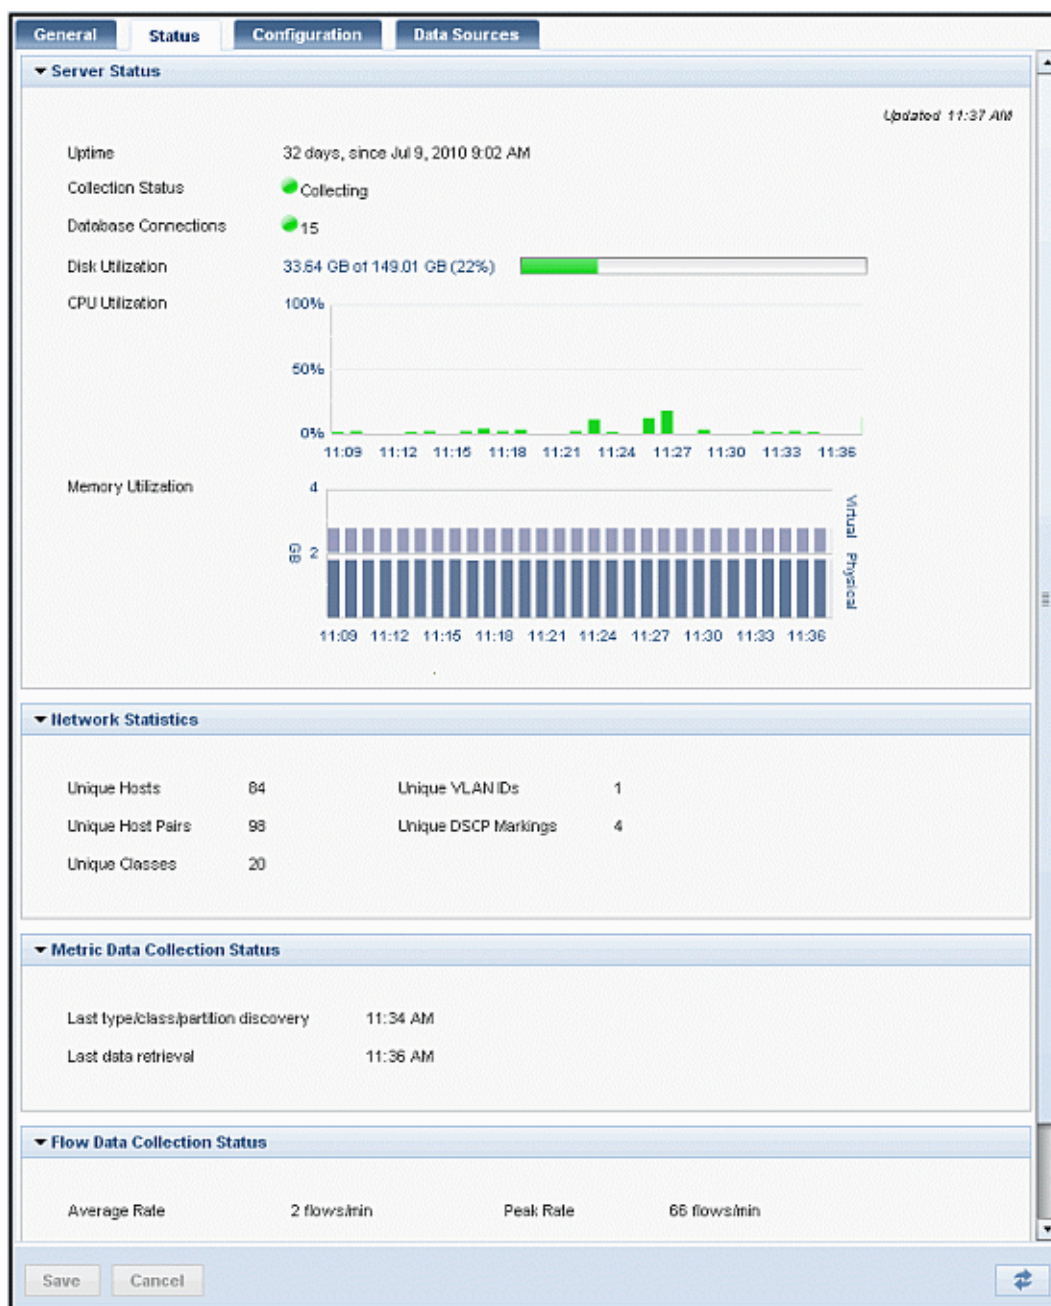
After you add DataCollector as a device, you can administer it from IntelligenceCenter. Other than [adding data sources](#), DataCollector does not require any additional configuration to function. However, you may find that you want to tune the default values to better align with your specific implementation. For example, if you do not generate any daily reports within your organization, you may want to reduce the amount of daily data that DataCollector stores in order to [conserve disk space](#).

To modify the data collection settings:

1. Select **Configure > Network > Devices**.

Note: To expand or collapse a network branch, click the arrow icons next to each branch. The down arrow ▼ icon indicates that the branch is expanded; the right arrow ► icon indicates that the branch is collapsed.

2. Select the DataCollector on the **Devices** tab. The DataCollector is indicated by the  icon. When you select the object, the DataCollector **Status** tab is displayed in the right-hand pane.



- Click the **Configuration** tab.

DataCollector1

General **Status** **Configuration** **Data Sources**

▼ **Metric Data Collection**

Query interval (minutes) 15 ▼

Data granularity (minutes) 5 ▼

Class discovery interval (hours) 24 ▼

Retention:

	Class	Partition	Link
Raw (hours)	48 ▲▼	48 ▲▼	48 ▲▼
Hour (hours)	48 ▲▼	48 ▲▼	48 ▲▼
Day (days)	32 ▲▼	32 ▲▼	32 ▲▼
Month (months)	7 ▲▼	7 ▲▼	7 ▲▼
Year (years)	2 ▲▼	2 ▲▼	2 ▲▼

▼ **Flow Data Collection**

Data port 9800

Raw (hours) 48 ▲▼

Hour (hours) 48 ▲▼

Day (days) 32 ▲▼


Month (months) 7 ▲▼

Year (years) 2 ▲▼

CSV ☒ On ☐ Off

CSV Duration (hours) 1 ▲▼

CSV Location C:\FlowRecords

Save **Cancel** 

4. If you want to customize how DataCollector interacts with data sources, modify the values in the *Data Collection Configuration* section:

Field	Description
Query interval (minutes)	The interval at which DataCollector polls an ME data source (default = 15 minutes). If you want to reduce the amount of traffic between DataCollector and the data sources, you can increase this interval. However, keep in mind that you will lose data granularity by doing so. Similarly, to increase data granularity (at the expense of more network traffic), reduce this value.
Data granularity (minutes)	The number (in minutes) that is used to divide the query interval into samples (default = 5 minutes). For example, if the query interval is 15 minutes and the data granularity is 5 minutes, the query results will be divided into 3 samples (15/5). To reduce the amount of aggregation that occurs on each sample of raw data, you can decrease the data granularity in relation to the query interval. Similarly, to increase the amount of summarization that is done on the raw data, you can increase the data granularity in relation to the query interval. Keep in mind that the more granular the data, the more disk space and processing power it requires.
Class discovery interval (hours)	The interval at which the DataCollector polls its data sources for class and partition information (default = 1 hour).
FDR data port	The UDP port on which the DataCollector listens for FDR data (default = 9800). The DataCollector must use the same port that your PacketShaper appliances and other network devices (such as routers that are emitting NetFlow-5 to DC) are configured to use for FDR.

5. To adjust the length of time that DataCollector stores ME data, modify the values in the *Metric Data Retention Configuration* section.

Note: Although DataCollector stores ME data for the specified amount of time, you should regularly [back up your IC and DC databases](#).

You can set different data retention values for each type of ME data (Class, Partition, and Link) as follows:

Field	Description
Raw (hours)	The number of hours of raw ME data that the DataCollector stores in its raw data table (default = 48 hours)
Hour (hours)	The number of hours worth of ME data that the DataCollector stores in its hourly table (default = 48 hours)
Day (days)	The number of days worth of ME data that the DataCollector stores in its daily table (default = 32 days)
Month (months)	The number of months worth of ME data that the DataCollector stores in its monthly table (default = 7 months)
Year (years)	The number of years worth of ME data that the DataCollector stores in its yearly table (default = 2 years)

6. If you want to modify the length of time that DataCollector stores Packeteer-2 or NetFlow-5 data, modify the values in the *Flow Data Retention Configuration* section. You can set different data retention values for each type of FDR data (Packeteer-2 or NetFlow-5) as follows:

Field	Description
Raw (hours)	The number of hours of raw FDR data that the DataCollector stores in its raw data table (default = 48 hours)
Hour (hours)	The number of hours worth of FDR data that the DataCollector stores in its hourly table (default = 48 hours)
Day (days)	The number of days worth of FDR data that the DataCollector stores in its daily table (default = 32 days)
Month (months)	The number of months worth of FDR data that the DataCollector stores in its monthly table (default = 7 months)

Year (years)	The number of years worth of data that the DataCollector stores in its yearly table (default = 2 years)
--------------	---

7. If you want to set up FDR data backup to a comma-separated values (CSV) file, modify the values in the *Flow Data Retention Configuration* section as follows:

Field	Description
CSV	Indicates whether you want to back up your raw flow data to a CSV file. Note that CSV export of FDR data supports only "Flow" tables (files prefixed with "f") and "VoIP" tables (files prefixed with "v"). Other tables are not backed up.
CSV Duration (hours)	The number of hours worth of raw flow data to back up into each CSV file (default = 1 hour). When the file reaches the specified duration, it is deleted and a new file starts.
CSV Location	The path to the directory on the DC server where you want to create the CSV files. Note that this directory must already exist.


8. When you are done modifying the settings, click **Save**.

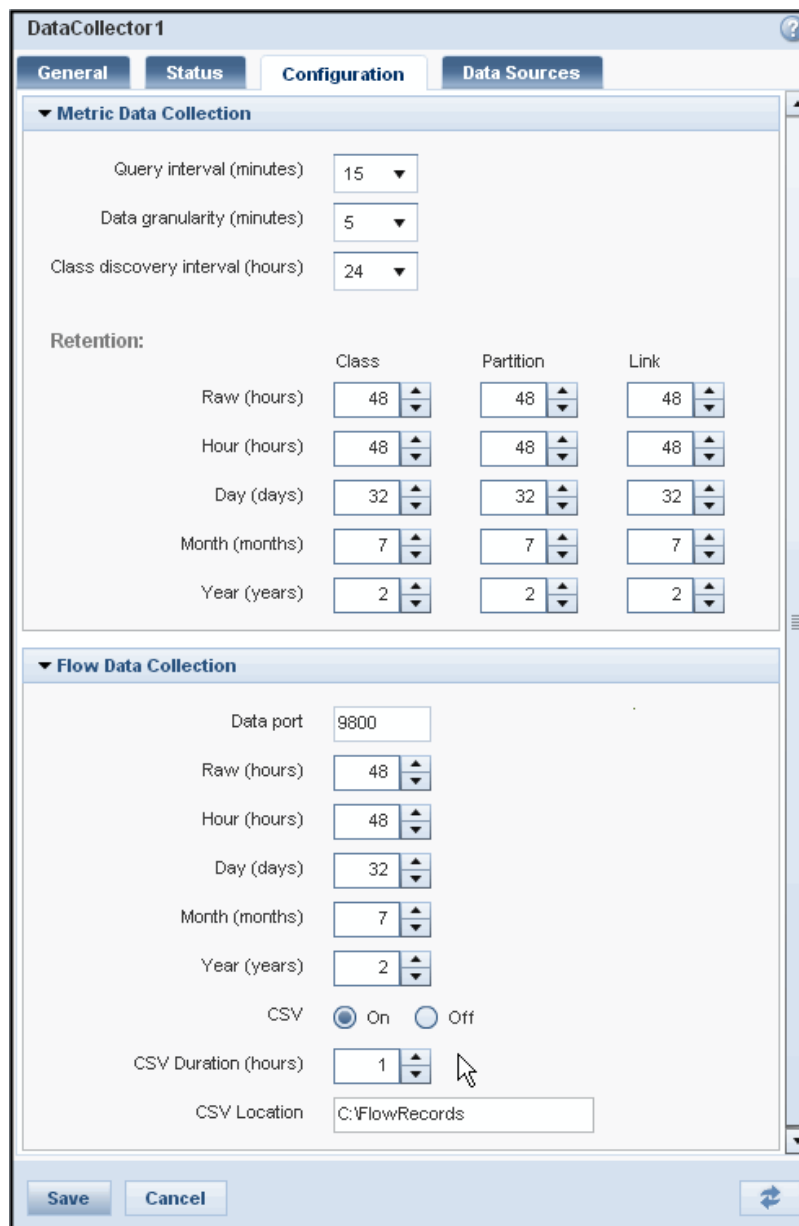
Back up FDR Data to a CSV File

To back up your [Packeteer-2](#) and or [NetFlow-5](#) data to a comma-separated values (CSV) file:

1. Select **Configure > Network > Devices**.

Note: To expand or collapse a network groups and sub-groups, click the arrow icons next to each group. The down arrow ▼ icon indicates that the group is expanded; the right arrow ► icon indicates that the group is collapsed.

2. Select the DataCollector object. The DataCollector is indicated by the  icon. When you select the object, the DataCollector **Status** tab is displayed in the right-hand pane.
3. Click the **Configuration** tab.



The screenshot shows the 'DataCollector1' configuration window with the 'Configuration' tab selected. The window is divided into two main sections: 'Metric Data Collection' and 'Flow Data Collection'.

Metric Data Collection:

- Query interval (minutes): 15
- Data granularity (minutes): 5
- Class discovery interval (hours): 24

Retention:

	Class	Partition	Link
Raw (hours)	48	48	48
Hour (hours)	48	48	48
Day (days)	32	32	32
Month (months)	7	7	7
Year (years)	2	2	2

Flow Data Collection:

- Data port: 9800
- Raw (hours): 48
- Hour (hours): 48
- Day (days): 32
- Month (months): 7
- Year (years): 2
- CSV: ☒ On ☐ Off
- CSV Duration (hours): 1
- CSV Location: C:\FlowRecords

At the bottom, there are 'Save' and 'Cancel' buttons, and a refresh icon.

4. To enable backup to a CSV file, set the value of the **CSV** field in the *Flow Data Collection* section to **On**.
5. Specify the number of hours worth of raw FDR data to back up into each CSV file (default = 1 hour) in the **CSV Duration** field. When the file reaches the specified duration, it is deleted and a new file starts.
6. Enter the path to the directory where you want the CSV file(s) saved in the **CSV Location** field. Note that this directory must already exist and that you cannot create it at the root of the drive.
7. Click **Save**.

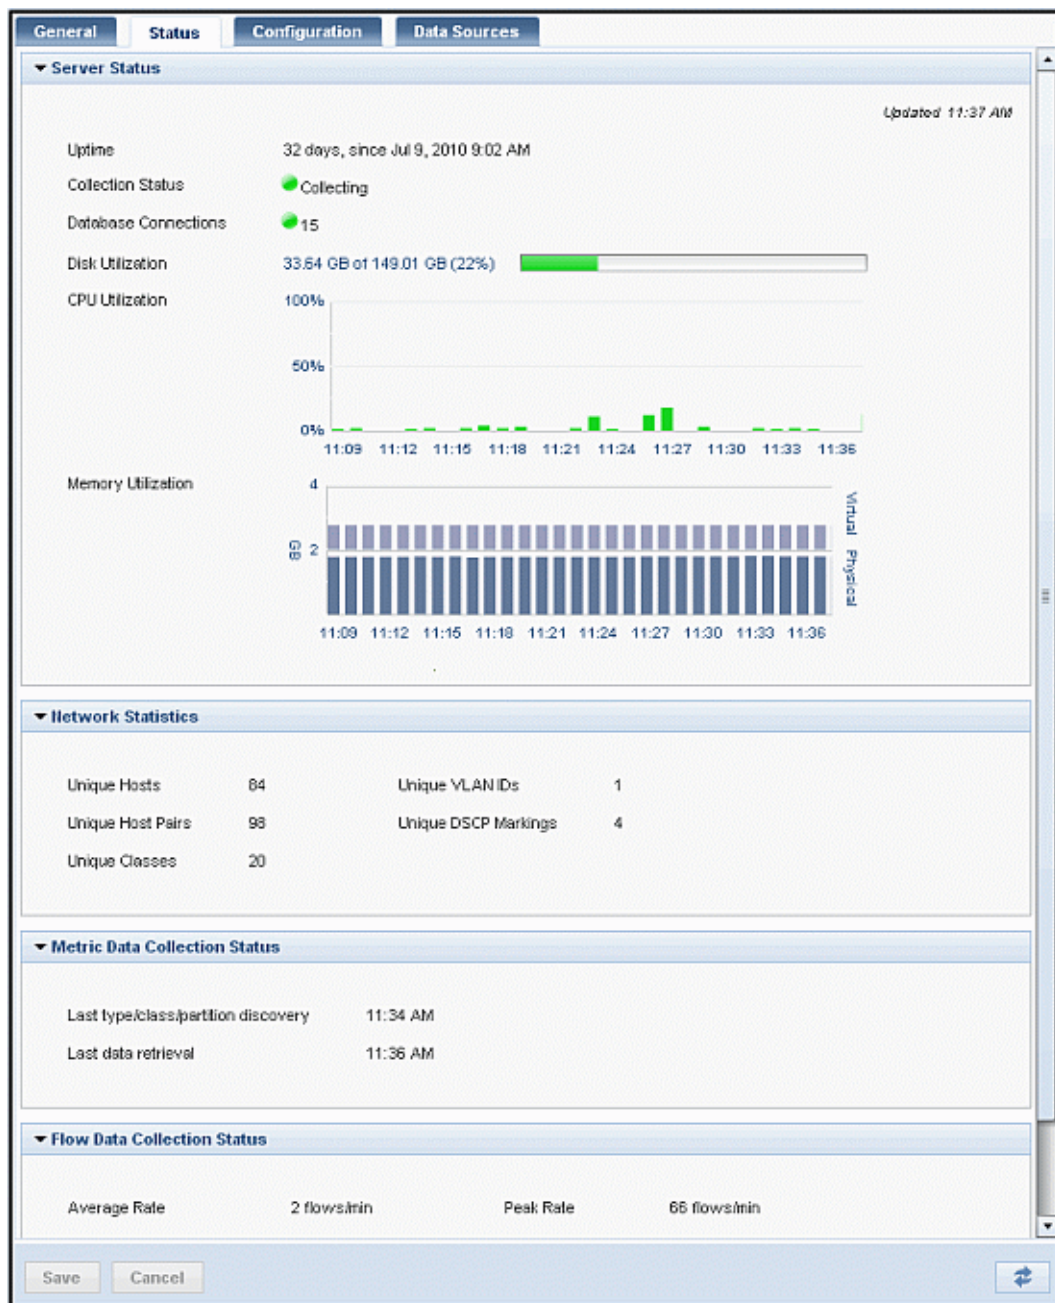
Remove a Data Source

To remove a data source from the DataCollector configuration:

1. Select **Configure > Network > Devices**.
2. Select the DataCollector object on the **Devices** tab. The DataCollector is indicated by the  icon.

Note: To expand or collapse a network group or sub-group, click the arrow icons next to each branch. The down arrow ▼ icon indicates that the group is expanded; the right arrow ► icon indicates that the group is collapsed.

When you select the object, the DataCollector [Status tab](#) is displayed.



3. Select the **Data Sources** tab. The data sources that you have configured are displayed.
4. To remove the data source, click the **Remove** link in the data source row. When prompted to confirm the delete operation, click **Yes**.

General Status Configuration Data Sources			
Source	Collect Metrics	Collect Flows	
▼ North America	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Remove
East Coast	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
France	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Remove

5. Click **Save**.

Monitor DataCollector Health


To monitor DataCollector:

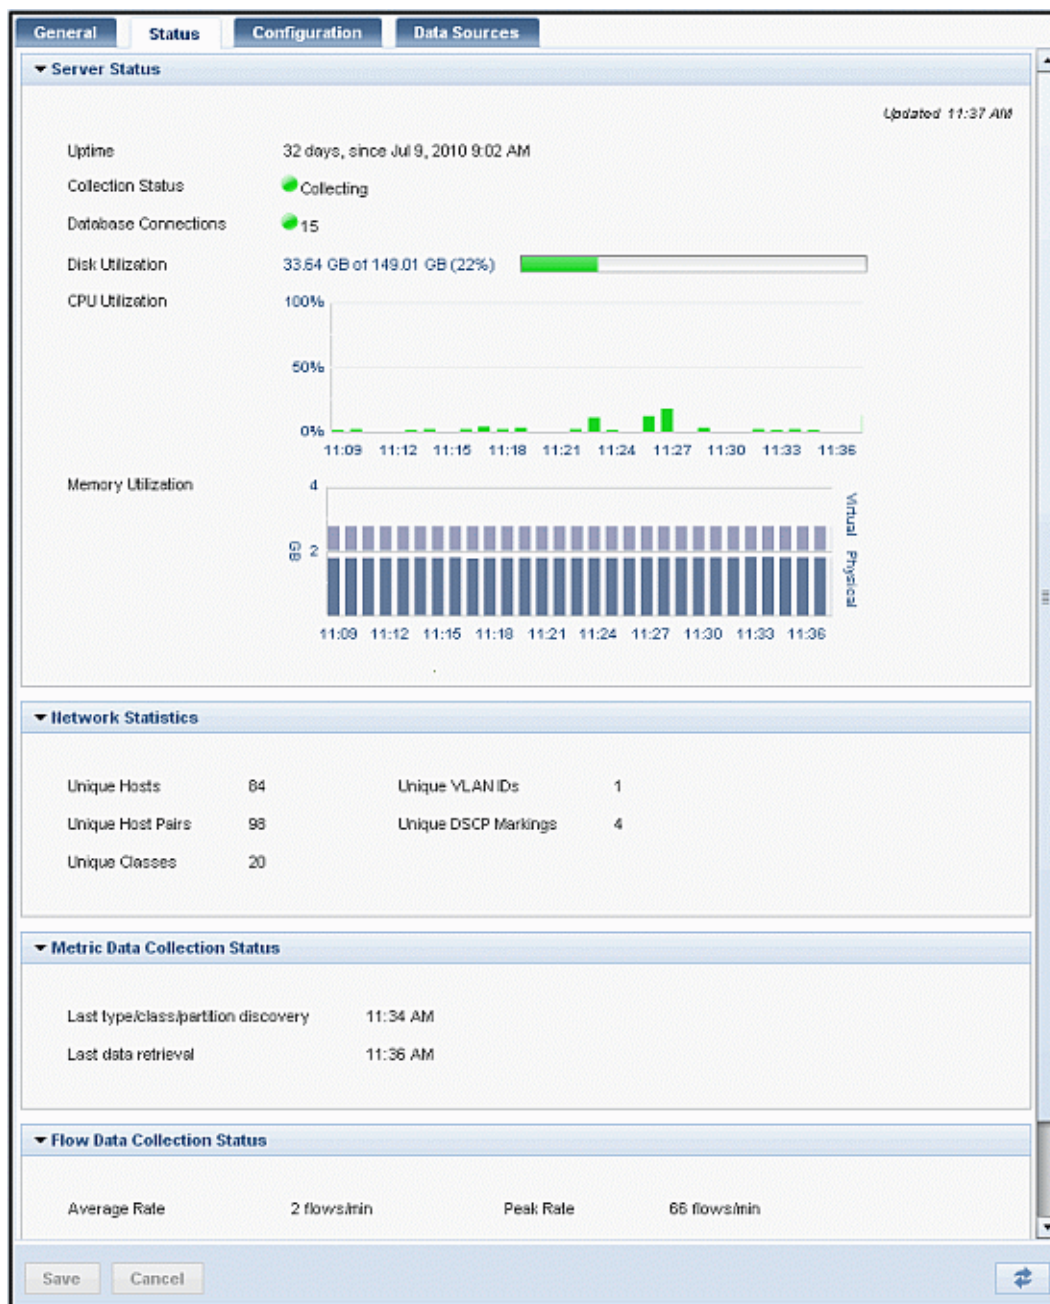
1. Select **Configure > Network > Devices**.
2. Select the DataCollector object on the **Devices** tab. The DataCollector is indicated by the  icon.

Note: To expand or collapse a network branch, click the arrow icons next to each branch. The down arrow ▼ icon indicates that the branch is expanded; the right arrow ► icon indicates that the branch is collapsed.

When you select the object, the DataCollector **Status** tab is displayed.





This screen displays health information about the server on which DataCollector is installed, status information about metric and flow data collection, and statistical information about the data that has been collected. In addition, if there are any collection errors, they are displayed at the bottom of the screen.

The status information on this screen is dynamically updated (every 60 seconds by default unless you've changed the status interval in [your IC profile](#)). You can also click the update  button in the lower right of the screen to force IntelligenceCenter to update the status now.



This tab displays the following information:

Field	Description
Server Status	
Uptime	The number of days since and the date and time that DataCollector last started.
Collection Status	Indicates whether data collection has been enabled. Possible

	<p>values are:</p> <ul style="list-style-type: none">  Collecting — DataCollector has started and is ready to collect data. Note that this does not indicate that DataCollector is actually collecting any data. See the <i>Metric Collection Status</i> and <i>Flow Data Collection Status</i> sections of this screen to determine whether data is being collected.  Stopped — DataCollector has stopped and is not collecting data. Possible reasons include license expiration, mismatched versions of IC and DC, or disk utilization reaching full capacity.  Migrating — DataCollector is migrating the database from a previous software release.  Initializing — DataCollector is starting up. You should only see this status for the first few minutes after a restart of the system.
Database Connections	The number of open connections to DataCollector's Postgres database. Connections get opened whenever DataCollector inserts data or whenever a user performs report or portlet operations from IntelligenceCenter.
Disk Utilization	The amount of disk space used by the DataCollector database. This field shows both the amount and percentage of disk space used and the total available disk space. It also shows a status bar that indicates whether you are in danger of running out of disk space. When the database reaches 95% of its disk space capacity, DC will stop collecting data and reporting will be interrupted. If you notice that disk utilization is nearing the 95% mark, you should take action to free up disk space on the drive before collection stops.
CPU Utilization	Shows CPU utilization over the last 10 minutes on the server where DataCollector is installed.
Memory Utilization	Shows virtual and physical memory utilization over the last 10 minutes on the server where DataCollector is installed.
Network Statistics <i>This section shows statistics about the flow data that DataCollector has collected. If you are not collecting FDR, these values will be 0.</i>	
Unique Hosts	The number of unique host entries in the DataCollector database.
Unique Host Pairs	The number of unique host pair entries in the DataCollector database.

Unique VLAN IDs	The number of unique VLAN IDs in the DataCollector database.
Unique Classes	The number of unique PacketShaper traffic classes in the DataCollector database.
Unique DSCP Markings	The number of unique Differentiated Services Code Point (DSCP) values in the DataCollector database.
Metric Data Collection Status <i>This section shows information about ME data collection and indicates if there are errors with your ME data sources.</i>	
Last link/class/partition discovery	Date and time that DataCollector last discovered traffic class and partition data from its data sources.
Last data retrieval	Date and time that DataCollector last collected ME data from one of its ME data sources.
x of y Configured data sources have errors	<p>This section is only displayed if one or more of your ME data sources has errors, where x indicates the number of configured ME data sources with errors and y indicates the total number of configured ME data sources.</p> <p>Each error lists the IP address of the Shaper that had the error and provides a brief description. To view ME data collection statistics for the device, click the right-arrow ► icon next to error message entry. You can also click the IP address to sign in to the appliance and conduct further troubleshooting.</p>
Flow Data Collection Status <i>This section shows information about FDR data collection and indicates if there are errors with your FDR data sources.</i>	
Average Rate	The average number of flows per minute received since DataCollector last started.
Peak Rate	The peak number of flows per minute received since DataCollector last started.
Invalid Records	The number of invalid FDR records that DataCollector received. Invalid records include records from configured data sources that are in a format not supported by DataCollector, such as intermediate FDR (iFDR).
Rejected Records	The number of FDR records that DataCollector rejected. DataCollector only accepts FDR records from its configured data sources .

Total Records	The total number of flow records collected.
Total Duplicate Records	Of the total flow records collected, the number that DC has determined to be duplicate. Duplicate records may be collected if the same flow travels through multiple PacketShaper appliances.
Flow Data Collection Errors	This section is only displayed if DC has encountered FDR errors. Each FDR error message shows the data and time of the FDR, the type of FDR (Packeteer-2 or NetFlow-5), the IP address of the device that sent the FDR, and a description of the error (for example, <i>Device not configured</i>).

Manage User Access

To access IC, a user must provide a valid user name and password. IC can authenticate users locally or it can contact an external authentication server to authenticate users. There are a couple of ways to create [user accounts](#):

- You can [manually create a local user account](#) for each IC user. This is a good method to use if you have few IC users and you do not have an existing authentication system in place.
- You can configure IC to authenticate users via an [external authentication](#) service (LDAP, RADIUS, or TACACS+). This is a good method to use if you already have authentication services in place and you do want to have to manually maintain additional user accounts. Keep in mind that for each user that is authenticated to IC via an external authentication service, you must manually [enable the account and assign a role](#) before the user can use IC.

Additionally, you must assign users to a role before they will be able to perform tasks within IC. You can [assign individual users to roles](#), or you can [create user groups](#) and assign roles to users via user group membership.

User Access Overview

To access and use IC, every user needs two things:

- **an enabled user account**—to log in to IC, a user must have a valid and enabled IC account. There are two ways to create a user account in IC: you can [manually create a user account](#) or IC can automatically create IC user accounts after authenticating the user credentials on an external authentication server. When IC automatically creates an account, the account does not get enabled. You must enable the account and assign a role to the user before the user will be able to log in to IC.
- **an assigned role**—The [role](#) or roles that are assigned to a user define the tasks the user can perform. Every user who uses IC must be assigned a role. The access rights that have been defined for the role determine what objects the users can see and what tasks the user can perform in IC; if the user doesn't have access rights to an object or a task, the user will not be able to see the corresponding menu item or object within the IC user interface. IC provides three default roles—*administrator*, *operator*, and *user*—that you can select from. Or, you can [define your own roles](#) and then assign them to users or groups of users.

After users successfully log in to IC, they can perform the tasks that their assigned roles allow. Additionally, the [single-sign on](#) credentials, [portlet settings](#), and [user profile preferences](#) that they configure (if they have rights to do so) are saved in the IC database.

Add a User Profile

A user profile consists of the user's contact information (including email address and phone number) as well as login information (user name and password). If you are using external authentication, a user profile will get created automatically whenever a user successfully authenticates. In this case, you will not need to create a user profile for the user; you will simply need to [activate the user's account](#).

If you are not using external authentication, you must add user profiles manually as follows:

1. Select **Access > Users & Groups > Users**. The user management window opens. The Users & Groups pane displays the existing user profiles. IntelligenceCenter provides one default user profile: *Administrator Administrator*. Configuration details for the selected user profile are displayed in the right-hand pane.
2. Click **Add**. The *Create User* dialog box is displayed.
3. Enter the **First Name**, **Last Name**, and, optionally, the **Email** address of the user.
4. Fill in the following user account information:

Field	Description
User Name	Unique name for the user — this is the name the user must enter in order to log in to IntelligenceCenter. The user name must be 32 characters or less; can contain any combination of letters, numbers, and symbols; and contain no spaces. User names are case sensitive. Note: The user name cannot be modified after adding the user.
New Password	The IntelligenceCenter password associated with the user name. The password must be 6 to 32 characters in length; can contain any combination of letters, numbers, and symbols; and contain no spaces. As you type the password, an asterisk (*) displays for each character you type. Passwords are case sensitive.
Verified Password	For verification purposes, this password must match the password you entered in the New Password field.

5. Click **Save**. The new user — identified by first and last name — appears in the *Users & Groups* list in the left pane. Configuration details for the user are displayed in the right-hand pane. By default, the **Account Status** is *Disabled*.

6. To set the user access role or account status, make sure the user object is selected in the *Users & Groups* pane the modify the following fields in the right pane:

Field	Description
Roles	<p>Select the access role you want this user to have. You can select one of the following default roles or one of the custom roles you created:</p> <p>user — Can view anything within IntelligenceCenter, but can not change anything. This access role is selected by default.</p> <p>operator — Can run and view reports and portlets and manage the IC network.</p> <p>administrator — Can perform all IC tasks. Administrators are the only users who can schedule reports, define report schedules, manage user profiles, upload license files, and modify the login message.</p>
User Groups	<p>You can add the user to any of the user groups that you have created by clicking the Add/Remove button in the <i>User Groups</i> area.</p>
Account Status	<p>If Enabled, the user can access IntelligenceCenter with the defined user name and password.</p> <p>If Disabled, the user cannot access IntelligenceCenter.</p>

7. **Note:** Required fields are marked with a red asterisk (*). If you do not complete a required field or if you enter an incorrect value, the field text box will be outlined in red. Users cannot access IntelligenceCenter unless the required **User Account Information** is filled in and the **Account Status** is *Enabled*.
8. Click **Save**.

Modify a User Profile

If you need to update a user's contact or login information or if you want to enable or disable a user's access to IC, you must modify the [user profile](#).

To modify a user profile:

1. Select **Access > Users & Groups**.
2. In the *Users & Groups* pane, select the user profile to modify. When you select a user, details about the user profile display in the right-hand pane.

User Name: * tandrew

Account Status: ☒ Enabled ☐ Disabled

New Password:

Verified Password:

First Name: * Timothy

Middle Name:

Last Name: * Andrew

Email: tandrew@acme.com

User Groups:

Roles:

Add Remove

3. Modify the fields as desired. Keep in mind that fields with a red asterisk (*) are required fields.
4. If you want to enable or disable the user profile, select the appropriate radio button in the **Account Status** field:
 - If **Enabled** is selected, the user can access IntelligenceCenter with the defined user name and password.
 - If **Disabled** is selected, the user cannot access IntelligenceCenter.
5. Click **Save**.

Note: Users are allowed to [change the password](#) associated with their own user profiles.

Manage Your User Profile

Your user profile contains settings specific to your IntelligenceCenter user sessions, including the screen that displays when you log in, the format used to select time settings, the interval at which dynamic information is updated on the screen, as well as your password. Your user profile password secures access to IntelligenceCenter. If you feel that your password has been compromised, you should change it.

Note: Only administrators can change another user's password.

To change user profile settings:

1. Click your username link in the IntelligenceCenter banner.



The *Configure User Profile* dialog box appears.

 A screenshot of the 'Configure User Profile' dialog box. It has a title bar with a question mark and a close button. The dialog is divided into two main sections: 'Change Password' and 'Preferences'.

The 'Change Password' section contains three text input fields: 'Current Password', 'New Password', and 'Verified New Password'. A 'Save' button is to the right of the 'New Password' field.

The 'Preferences' section contains three settings:

- 'Time Entry Format' with two radio buttons: '12 hours' (selected) and '24 hours'.
- 'Initial Start-Up Screen' with a dropdown menu showing 'Monitor' and a 'Save' button to its right.
- 'Device Refresh Interval' with a spinner box showing '1' and the unit 'min'.

 A 'Close' button is located at the bottom right of the dialog.

2. To change your password:
 - Type your **Current Password**.
 - Type your **New Password**. The password must be from 6 to 32 characters in length and can contain any combination of letters, numbers, and symbols and must not contain any spaces. As you type the password, an asterisk (*) displays for each character you type. Passwords are case sensitive.
 - Retype the password in the **Verified New Password** field.
 - Click **Save**.
3. To set preferences for the IntelligenceCenter application:
 - If you want to set the time format that is displayed on time entry screens, select a **Time Entry Format** radio button. By default, time settings are displayed using a 12-hour clock format.

- If you want to change the screen that is displayed when you first log in to IntelligenceCenter, select a tab from the **Initial Start-Up Screen** drop-down list. By default, the **Monitor** tab is displayed. However, you can choose any of the following as your initial start-up screen: **Monitor**, **Run Reports**, **View Reports**, **Network**, **Applications**, **Schedules**, **View Alerts**, **Configure Alerting**, **Users & Groups**, **Roles**, **Audit**, **License**, **Scheduled Tasks**, or **System Settings**.
 - If you want to change the interval at which IC polls DataCollector (DC) for updated status information, enter a new value in the **Device Refresh Interval** field. By default, IC updates DC status information every minute; the range is 1 minute to 12 minutes. Keep in mind that the lower the value, the higher the network traffic and CPU usage.
 - Click **Save**.
4. When you finish updating your user profile, click **Close**.

Delete a User Profile

When a user profile is no longer required, it can be deleted. If you only want to temporarily prevent a user from accessing IntelligenceCenter, don't delete the account — disable the user access. (See [Modify a User Profile](#).)

To delete a user profile:

1. Select **Access > Users & Groups**.
2. In the *Users* pane, select the user to delete.
3. Click **Remove**.
4. Confirm the deletion.

Note: The Administrator Administrator user cannot be deleted.

Configure User Groups

You must assign users to a role before they will be able to perform tasks within IC. You can [assign individual users to roles](#), or you can [create user groups](#) and assign roles to users via user group membership.

Add a User Group

To create a user group:

1. Select **Access > Users & Groups** and select the **User Groups** tab. The *Users & Groups* pane displays the existing user groups. There are no default user groups.
2. Click **Add** to create a new user group. The *Create User Group* dialog box is displayed.
3. Enter a **User Group Name** and optionally a **Description** of the group and then click **Save**. The new user group is added to the **User Groups** tab and details about the group are displayed in the right pane.

The screenshot shows the 'Users & Groups' application. The left pane has a 'Users & Groups' header with 'Users' and 'User Groups' tabs. The 'User Groups' tab is selected, showing a list of user groups. The 'Application Admins' group is highlighted. The right pane, titled 'Application Admins', displays the configuration for this group. It includes a 'User Group Name' field with the value 'Application Admins', a 'Description' field, a 'Roles' section with a note: 'Note: Assigning roles will not make Users with these roles members of this User Group. Individual Users and other User Groups will need to be added to the member list in order to be a part of this User Group.', and a 'Members' section. Each section has an 'Add/Remove' button.

4. To add roles that you want to assign to the members of this group, click the **Add/Remove** button in the **Roles** section of the screen. The *Add/Remove Roles* dialog box is displayed.
5. Select the role(s) you want to associate with this user group and then click **Make Changes**. The selected roles are displayed in the **Roles** box.
6. To add users to this user group, click the **Add/Remove** button in the **Members** section of the screen. The *Add/Remove Members* dialog box is displayed.
7. Select the users and user groups you want to associate with this user group and then click **Make Changes**. The selected users and user groups are displayed in the **Members** box.
8. To save the user group configuration, click **Save**.

Modify a User Group

To modify a user group:

1. Select **Access > Users & Groups** and select the **User Groups** tab. The *Users & Groups* pane displays the existing user groups.
2. Select the user group you want to modify. Details about the group are displayed in the right pane.

The screenshot shows the 'Application Admins' configuration interface. On the left, a sidebar titled 'Users & Groups' has tabs for 'Users' and 'User Groups'. Under 'User Groups', 'Application Admins' is selected. The main area is titled 'Application Admins' and contains four sections: 'User Group Name' (with a red asterisk and the text 'Application Admins'), 'Description' (an empty text box), 'Roles' (with a note: 'Note: Assigning roles will not make Users with these roles members of this User Group. Individual Users and other User Groups will need to be added to the member list in order to be a part of this User Group.' and an empty list box), and 'Members' (an empty list box). Each of the 'Roles' and 'Members' sections has an 'Add/Remove' button.

3. To add roles that you want to assign to the members of this group, click the **Add/Remove** button in the **Roles** section of the screen. The *Add/Remove Roles* dialog box is displayed.
4. Select the role(s) you want to associate with this user group and then click **Make Changes**. The selected roles are displayed in the **Roles** box.
5. To add users to this user group, click the **Add/Remove** button in the **Members** section of the screen. The *Add/Remove Members* dialog box is displayed.
6. Select the users and user groups you want to associate with this user group and then click **Make Changes**. The selected users and user groups are displayed in the **Members** box.
7. To save the user group configuration, click **Save**.

Define Group Membership

When you add a [user](#) as a member of a group, the user is granted the access levels defined for all roles that are assigned to the user group.

To define group membership:

1. Select **Access > Users**.
2. There are two ways to add members to a user group:
 - To define group membership for an individual user, select the **Users** tab and then select the user you want to add to a group. The user account details are displayed in the right pane. Click the **Add/Remove** button in the **User Groups** section of the screen. The *Add/Remove User Groups* dialog box is displayed. Select the **User Groups** you want to add the user to and then click **Make Changes**.
 - To define group membership for multiple users, select the **User Groups** tab and then select the user group to which you want to add users. The user group details are displayed in the right pane. Click the **Add/Remove** button in the **Members** section of the screen. The *Add/Remove User Members* dialog box is

displayed. Select the **Users** and **User Groups** you want to add the user to and then click **Make Changes**.

3. Click **Save**.

Enable External Authentication



As an alternative to [configuring a local user](#) account for each IC user, you can configure IC to authenticate users by accessing your existing RADIUS, TACACS+, and/or LDAP authentication services. If external authentication is enabled, IC will contact the configured authentication server(s) to authenticate a user based on the user name and password supplied during login. The first time IC successfully authenticates a user using an external authentication service, it creates an IC [user account](#) for the user. However, IC will not log the user in upon this first successful authentication. Instead, it displays a message on the login screen indicating that the administrator must activate the account; in addition, IC generates an [alert](#) and sends an email to the IC [system email account](#) indicating that the account requires activation.

The following sections describe how to configure IC to use one or more external authentication services:

- [Add an external authentication service](#)
- [Configure RADIUS](#)
- [Configure TACACS+](#)
- [Configure LDAP](#)
- [Activate a user account](#)

Add an External Authentication Service

To configure IC to work with an external authentication server::

1. Select **Manage > System Settings > External Authentication**.
2. In the *External Authentication* pane, click **Add**. The *Add an External Authentication Service* dialog box is displayed.
3. Select the type of authentication service you are adding from the **Please select the service you want to add** drop-down list. When you make a selection, the configuration fields that pertain to the selected authentication service are displayed on the dialog box.
4. Enter the configuration information that IC needs to connect to the external authentication server. The configuration information you must supply depends on which service you selected:
 - [Configure RADIUS](#)
 - [Configure TACACS+](#)
 - [Configure LDAP](#)
5. Click **Save**. The authentication service is added to the **Services** table in the *External Authentication* pane. If you had already added a different external authentication service, the new service is added to the bottom of the list. When a user attempts to authenticate, IC will connect to the configured authentication services in the order listed.
6. If you want to change the order of the authentication services, select the service you want to move in the **Services** list and click the up arrow  or the down arrow  to move the service to the desired position and then click **Save**.

Configure RADIUS

To configure IC to work with the RADIUS authentication server:

1. Select **Manage > System Settings > External Authentication**.
2. In the *External Authentication* pane, click **Add**. The *Add an External Authentication Service* dialog box is displayed.
3. Select **RADIUS** from the **Please select the service you want to add** drop-down list. The RADIUS configuration fields are displayed on the dialog box.
4. Enter the following configuration settings for the primary and optionally for the secondary RADIUS authentication server:

Field	Description
Host (Primary or Secondary)	The IP address or DNS name of the RADIUS server.
Port	The port number IC will use to connect to the RADIUS server. The default port number is 1812.
Shared Secret	The password required for the RADIUS server to authenticate the IC server. For verification purposes, you must retype this password in the Retype Shared Secret field.
Authentication Scheme	<p>The method to be used to authenticate the user credentials. You can select one of the following methods:</p> <ul style="list-style-type: none"> • PAP—With PAP (Password Authentication Protocol), the user name and password are transmitted in clear, unencrypted text. If you select PAP, Blue Coat recommends you increase security by logging in to IC via HTTPS. PAP is required for RADIUS configurations that require access to clear text passwords (for example, when passwords are stored and maintained in a database external to the RADIUS server). • CHAP—With CHAP (Challenge Handshake Authentication Protocol) the RADIUS server sends a challenge that consists of a session ID and an arbitrary challenge string, and the user name and password are encrypted before they are sent back to the server.

4. You can optionally modify the **Retry Settings** as follows:
 - **Retry limit**—By default, if the RADIUS server fails to respond, IC will try to log in to the server three times before reporting a server failure. The valid range for this field is 0-10.
 - **Retry interval**—By default, the IC server waits 5 seconds before retrying a login when the RADIUS server fails to respond. The valid range for this field is 0-30 seconds.
5. Make sure the **Enable RADIUS** checkbox is checked.

- Click **Save**. The **RADIUS** service is added to the **Services** table in the *External Authentication* pane. If you had already added a different external authentication service, **RADIUS** is added to the bottom of the list. When a user attempts to authenticate, IC will connect to the configured authentication services in the order listed.
- If you want to change the order of the authentication services, select the service you want to move in the **Services** list and click the up arrow or the down arrow to move the service to the desired position and then click **Save**.

Configure TACACS+

To configure IC to work with the TACACS+ authentication server:

- Select **Manage > System Settings > External Authentication**.
- In the *External Authentication* pane, click **Add**. The *Add an External Authentication Service* dialog box is displayed.
- Select **TACACS+** from the **Please select the service you want to add** drop-down list. The TACACS+ configuration fields are displayed on the dialog box.
- Enter the following configuration settings for the primary and optionally for the secondary TACACS+ authentication server:

Field	Description
Host (Primary or Secondary)	The IP address or DNS name of the TACACS+ server.
Port	The port number IC will use to connect to the TACACS+ server. The default port number is 49.
Shared Secret	The password required for the TACACS+ server to authenticate the IC server. For verification purposes, you must retype this password in the Retype Shared Secret field.
Authentication Scheme	<p>The method to be used to authenticate the user credentials. You can select one of the following methods:</p> <ul style="list-style-type: none"> PAP—With PAP (Password Authentication Protocol), the user name and password are transmitted in clear, unencrypted text. If you select PAP, Blue Coat recommends you increase security by logging in to IC via HTTPS. PAP is required for TACACS+ configurations that require access to clear text passwords (for example, when passwords are stored and maintained in a database external to the TACACS+ server). CHAP—With CHAP (Challenge Handshake Authentication Protocol) the TACACS+ server sends a challenge that consists of a session ID and an arbitrary challenge string, and the user name and password are encrypted before they are sent back to the server. This is the default. ASCII—With ASCII (American Standard Code for

	Information Interchange), the user name and password are transmitted in clear, unencrypted text.
--	--

4. You can optionally modify the **Retry Settings** as follows:
5. **Retry limit**—By default, if the TACACS+ server fails to respond, IC will try to log in to the server three times before reporting a server failure. The valid range for this field is 0-10.
6. **Retry interval**—By default, the IC server waits 5 seconds before retrying a login when the TACACS+ server fails to respond. The valid range for this field is 0-30 seconds.
7. Make sure the **Enable TACACS+** checkbox is checked.
8. Click **Save**. The **TACACS+** service is added to the **Services** table in the External Authentication pane. If you had already added a different external authentication service, **TACACS+** is added to the bottom of the list. When a user attempts to authenticate, IC will connect to the configured authentication services in the order listed.
9. If you want to change the order of the authentication services, select the service you want to move in the **Services** list and click the up arrow or the down arrow to move the service to the desired position and then click **Save**.

Configure LDAP

To configure IC to work with the LDAP authentication server:

1. Select **Manage > System Settings > External Authentication**.
2. In the *External Authentication* pane, click **Add**. The *Add an External Authentication Service* dialog box is displayed.
3. Select LDAP from the **Please select the service you want to add** drop-down list. The LDAP configuration fields are displayed on the dialog box.
4. Enter the following configuration settings for the primary and optionally for the secondary LDAP authentication server:

Field	Description
Host (Primary or Secondary)	The IP address or DNS name of the LDAP server.
Port	The port number IC will use to connect to the LDAP server. The default port number is 389. If you enable SSL, change the port to an SSL listening port, such as port 636 (the default LDAPS port).
Type of server	Select the type of LDAP authentication server you are using: Microsoft Active Directory , Netscape/Sun iPlanet , Novell NDS/eDirectory , or Other (the default).
User Attribute Type	Specifies which LDAP attribute contains the user name. IC uses the following default settings, however you can change these values if your LDAP authentication server uses a different user attribute: <ul style="list-style-type: none"> • Microsoft Active Directory: sAMAccountName • Netscape/Sun iPlanet: uid

	<ul style="list-style-type: none"> • Novell NDS/eDirectory/Other: <code>cn</code>
LDAP Version	<p>Select 2 for LDAP v2 or 3 for LDAP v3 (the default). If you use LDAP v3, you can also enable the following options:</p> <p>Follow Referrals—Select this option if your LDAP authentication service comprises multiple servers (or domains) and you want IC to automatically be redirected to the server containing the information for the user it is trying to authenticate.</p> <p>Enable SSL—Select this option if you want IC and the LDAP server to communicate over a secure connection. Note that you will also need to change the Port to an SSL listening port, such as port 636 (the default LDAPS port).</p>
Base DNs	<p>A Base DN identifies the LDAP entry that IC should use as the starting point of its search for users. You can specify multiple base DNs. You must enter a complete DN, for example <code>cn=administrators, dc=acme, dc=com</code> or <code>ou=operations, o=acme</code>.</p>

4. You can optionally modify **Retry limit**. By default, if the LDAP server fails to respond, IC will try to log in to the server three times before reporting a server failure. The valid range for this field is 0-10.
5. Make sure the **Enable LDAP** checkbox is checked.
6. Click **Save**. The **LDAP** service is added to the **Services** table in the *External Authentication* pane. If you had already added a different external authentication service, **LDAP** is added to the bottom of the list. When a user attempts to authenticate, IC will connect to the configured authentication services in the order listed.
7. If you want to change the order of the authentication services, select the service you want to move in the **Services** list and click the up arrow or the down arrow to move the service to the desired position and then click **Save**.

Activate a User Account

The first time IC successfully authenticates a user using an external authentication service, it creates an [IC user account](#) for the user. However, IC will not log the user in upon this first successful authentication. Instead, it displays a message on the login screen indicating that the administrator must activate the account; in addition, IC generates an alert and sends an email to the IC [system email account](#) indicating that the account requires activation.

Before the externally authenticated users can access IC, you must activate their accounts and assign them to at least one role (or add them to a user group that is associated with a role) as follows:

1. Select **Access > Users & Groups** and select the **Users** tab. The *Users* pane displays the existing user accounts.
2. Select the user account that you want to activate. Details about the selected user account are displayed in the right pane.
3. In the **Account Status** field, select **Enabled**.
4. Assign at least one role to the user by selecting User Groups and/or Roles:

- If you want to associate the user with a user group that has roles assigned, click **Add/Remove** in the **User Groups** section of the screen, select the user groups you want to add the user to from the dialog box and then click **Make Changes**.
- If you want to manually assign the user to a role, click **Add/Remove** in the **Roles** section of the screen, select the roles you want to assign the user to from the dialog box and then click **Make Changes**.

5. Click **Save**.

Manage Roles

In IC, an access role defines the tasks a given user can perform. Every user who uses IC must be assigned a role. The access rights that have been defined for the role determine what objects the users can see and what tasks the user can perform in IC.

This section describes how to manage IC [roles](#). It includes the following topics:

- [Create a new role](#)
- [Define access levels for a role](#)
- [Add users to a role](#)
- [Delete a role](#)

Add a Role

To define a new role:

1. Select **Access > Roles**.
2. Click **Add/Copy**. The *Add a New Role* dialog box is displayed.

Add a New Role ✕

Select a radio button below to specify whether you want to create the new role from scratch or copy all attributes of an existing role to the new role:

☒ Create a New Role

Name this New Role: *

Description:

☐ Copy an Existing Role

Select a Role to Copy: * administrator ▼

Name this Copied Role: *

Description:

Add New Role

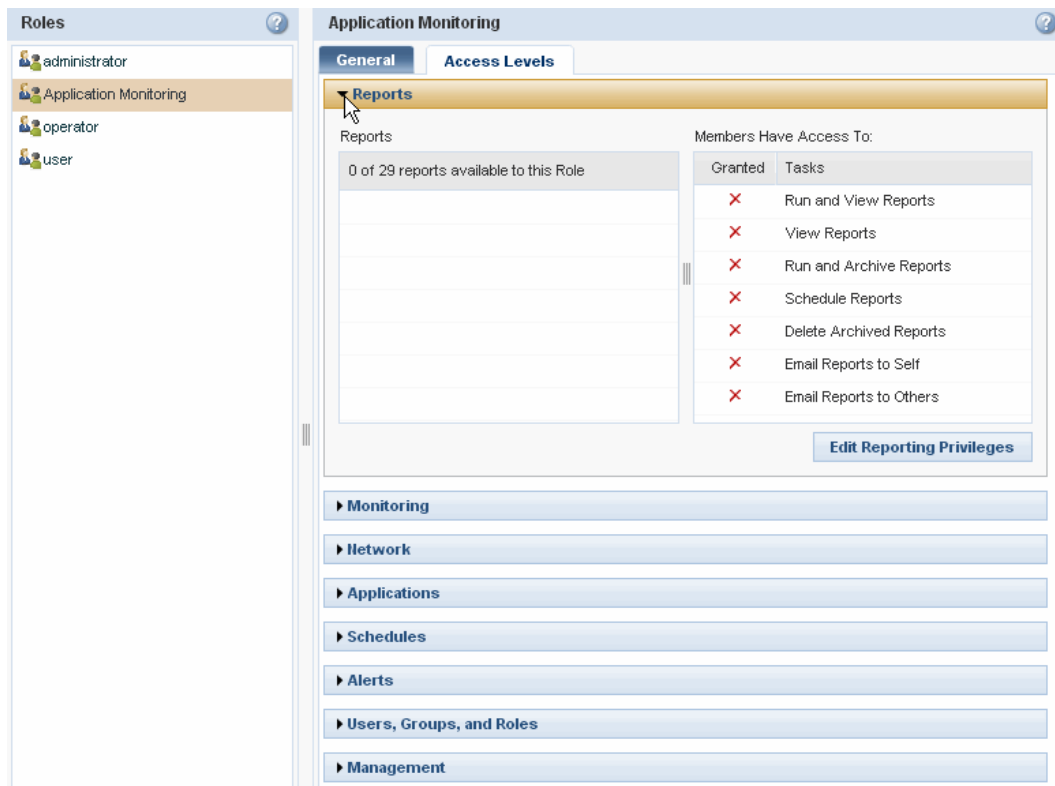
3. Specify how you want to create the role:
 - If you want to create a role and manually define the access privileges for the role, make sure the **Create a New Role** radio button is selected. When you create a role manually, all access privileges for all IC tasks are disabled by default.
 - If you want to create a new role that inherits the access privileges defined for an existing role, select the **Copy an Existing Role** radio button and then choose a role from the **Select a Role to Copy** drop-down list. You can modify the access privileges after you create the role.
4. Enter a name for the new role in the **Name this New Role** or **Name this Copied Role** field.
5. (Optional) Enter a **Description** for the new role.
6. Click **Add New Role**. The new role is added to the *Roles* pane. You can now [define the access levels](#) for the role and [add members](#).

Define Access Levels for a Role

The access rights that have been defined for a role determine what objects the members of the role can see and what tasks they can perform in IC. If a user doesn't have access rights to an object or a task, the user will not be able to see the corresponding menu item or object within the IC user interface. Therefore, when defining access levels for a role, keep in mind that there are [dependencies between access levels](#). For example, a user will not be able to modify an object if they do not have the right to view that object in the first place.

To define the access levels for the role:

1. Select **Access > Roles**.
2. Select the role you want to configure from the *Roles* pane (or [create a new role](#)).
3. In the right pane, select the **Access Levels** tab. This tab details the access levels for the various categories of tasks that users can perform from within IC.
4. Expand the category for which you want to define access levels by clicking the right arrow ► icon. After you expand the category, the current access levels for the category are displayed.



5. To set the access privileges for the category, click the corresponding **Edit** button. For example, to edit the access levels for reporting, click the **Edit Reporting Privileges** button. The privileges editor for the category is displayed. You use the privileges editor to define the access levels for the category, including which tasks the user can perform within the category and which objects within the category the user can perform these tasks on. The specific objects and/or tasks that are available depend on which category you selected.

Edit Reporting Privileges for Role Application Monitoring

Select the report categories and/or specific reports that you want users in this role to be able to access:

Reports View: All

- ▼ All
 - ▼ Application
 - ✓ Application Response Time
 - ✓ Application Activity
 - ✓ Top Immediate Children
 - ✓ Top Applications
 - ▶ Device
 - ▶ Host
 - ▶ Site

Select the reporting tasks that users can perform for the selected reports:

Granted

- ☐ Run and View Reports
- ☐ View Reports
- ☐ Run and Archive Reports
- ☐ Schedule Reports
- ☐ Delete Archived Reports
- ☐ Email Reports to Self
- ☐ Email Reports to Others

Save Changes

For example, in the reporting privileges editor (shown above), you can grant users in the role the privilege to perform a reporting task—such as viewing reports, emailing reports, or scheduling reports—by checking the corresponding checkbox in the **Granted** column. You can specify which reports the users in the role can perform these tasks on by selecting the specific reports in the **Reports** column. Many of the privileges editors have hierarchical lists of objects that you can expand and collapse by clicking arrow icons next to each group of objects. The down arrow icon indicates that the group is expanded; the right arrow icon indicates that the group is collapsed. Selecting an object at a higher level of the hierarchy automatically selects all objects below it. For example, selecting the **Application** object in the list of reports automatically selects all application reports. Similarly, by granting access to a specific network view, rights to see the view members are also granted.

6. When you are done defining the access privileges for the category, click **Save Changes**.
7. Repeat Step 4 through Step 6 for each category of the IC application for which you want to define access privileges for this role.

Access Levels for the Default Roles

IC provides three default roles—*administrator*, *operator*, and *user*. You can assign these default roles to your IC users, use these default roles as templates when creating new roles, or you can choose not to use these default roles and [create your own roles](#) from scratch. The default user *admin* is assigned the *administrator* role; you cannot delete this user.

For detailed information about which tasks each of the default roles can perform, select **Access > Roles > Access Levels**.

Access Level Dependencies

When defining access levels for a role, keep in mind that some rights are dependent on other rights. For example, suppose you wanted to create an application monitoring role for users who need to use IC to run application reports. To run an application report, the user needs to be able to select the report to run, select the network group, sub-group, or device against which to run the report, and, in some cases, select the actual application to report on. Therefore, when defining access levels for this role you would need to grant the following access levels to the role:

- **Reports**—Grant rights view the application reports and to perform the desired reporting tasks such as running, viewing, scheduling, and emailing reports.

Edit Reporting Privileges for Role Application Monitoring																	
<p>Select the report categories and/or specific reports that you want users in this role to be able to access:</p> <div> <div>Reports</div> <div>View: All</div> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> All <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Application <input type="checkbox"/> Device <input type="checkbox"/> Host <input type="checkbox"/> Site </div>	<p>Select the reporting tasks that users can perform for the selected reports:</p> <table border="1"> <thead> <tr> <th>Granted</th> <th></th> </tr> </thead> <tbody> <tr><td><input checked="" type="checkbox"/></td><td>Run and View Reports</td></tr> <tr><td><input checked="" type="checkbox"/></td><td>View Reports</td></tr> <tr><td><input type="checkbox"/></td><td>Run and Archive Reports</td></tr> <tr><td><input type="checkbox"/></td><td>Schedule Reports</td></tr> <tr><td><input type="checkbox"/></td><td>Delete Archived Reports</td></tr> <tr><td><input checked="" type="checkbox"/></td><td>Email Reports to Self</td></tr> <tr><td><input type="checkbox"/></td><td>Email Reports to Others</td></tr> </tbody> </table>	Granted		<input checked="" type="checkbox"/>	Run and View Reports	<input checked="" type="checkbox"/>	View Reports	<input type="checkbox"/>	Run and Archive Reports	<input type="checkbox"/>	Schedule Reports	<input type="checkbox"/>	Delete Archived Reports	<input checked="" type="checkbox"/>	Email Reports to Self	<input type="checkbox"/>	Email Reports to Others
Granted																	
<input checked="" type="checkbox"/>	Run and View Reports																
<input checked="" type="checkbox"/>	View Reports																
<input type="checkbox"/>	Run and Archive Reports																
<input type="checkbox"/>	Schedule Reports																
<input type="checkbox"/>	Delete Archived Reports																
<input checked="" type="checkbox"/>	Email Reports to Self																
<input type="checkbox"/>	Email Reports to Others																

- **Network > Devices**—Grant rights to the specific groups, sub-groups, and devices you want the users in the role to be able to report against and grant rights to view network groups, sub-groups, and devices.

Edit Device Privileges for Role Application Monitoring																	
<p>Select the groups, sub-groups, and devices that you want users in this role to be able to access:</p> <div> <div>Network Topology</div> <ul style="list-style-type: none"> <input type="checkbox"/> All <ul style="list-style-type: none"> <input type="checkbox"/> Network <ul style="list-style-type: none"> <input type="checkbox"/> Asia <input type="checkbox"/> DataCollector1 <input type="checkbox"/> Europe <input checked="" type="checkbox"/> North America <input type="checkbox"/> South America </div>	<p>Select the tasks that users can perform on the selected groups, sub-groups, and devices:</p> <table border="1"> <thead> <tr> <th>Granted</th> <th></th> </tr> </thead> <tbody> <tr><td><input type="checkbox"/></td><td>Modify Network Group, Sub-group, and Device Properties</td></tr> <tr><td><input type="checkbox"/></td><td>Delete Network Groups, Sub-groups, and Devices</td></tr> <tr><td><input checked="" type="checkbox"/></td><td>View Network Groups, Sub-groups, and Devices</td></tr> <tr><td><input type="checkbox"/></td><td>View Device Details (including Single Sign-On Credentials)</td></tr> <tr><td><input type="checkbox"/></td><td>Create Groups</td></tr> <tr><td><input type="checkbox"/></td><td>Create Sub-groups</td></tr> <tr><td><input type="checkbox"/></td><td>Create Devices (manual or by import)</td></tr> </tbody> </table>	Granted		<input type="checkbox"/>	Modify Network Group, Sub-group, and Device Properties	<input type="checkbox"/>	Delete Network Groups, Sub-groups, and Devices	<input checked="" type="checkbox"/>	View Network Groups, Sub-groups, and Devices	<input type="checkbox"/>	View Device Details (including Single Sign-On Credentials)	<input type="checkbox"/>	Create Groups	<input type="checkbox"/>	Create Sub-groups	<input type="checkbox"/>	Create Devices (manual or by import)
Granted																	
<input type="checkbox"/>	Modify Network Group, Sub-group, and Device Properties																
<input type="checkbox"/>	Delete Network Groups, Sub-groups, and Devices																
<input checked="" type="checkbox"/>	View Network Groups, Sub-groups, and Devices																
<input type="checkbox"/>	View Device Details (including Single Sign-On Credentials)																
<input type="checkbox"/>	Create Groups																
<input type="checkbox"/>	Create Sub-groups																
<input type="checkbox"/>	Create Devices (manual or by import)																

- **Applications**—Grant rights to view application definitions. You must also specify which applications users in the role can see or select All applications. Keep in mind that limiting the applications that the user can see will only limit what applications users can choose to run reports against; it does not limit what applications show up on the reports. For example, if you limit the application access to HTTP, FTP, and CIFS, users will only be able to run an Application Activity report against one of those applications. If, however, they run a Top Applications report the report will display the applications that used the most bandwidth during the reporting period, whether or not the users in the role have specifically been granted rights to view those applications.

Edit Application Privileges for Role Application Monitoring											
Select the applications that you want users in this role to be able to access:	Select the application tasks that users can perform on the selected applications:										
<div>Applications</div> <div> <input checked="" type="checkbox"/> All </div>	<table border="1"> <thead> <tr> <th>Granted</th> <th></th> </tr> </thead> <tbody> <tr> <td><input type="checkbox"/></td> <td>Create Application Definitions</td> </tr> <tr> <td><input type="checkbox"/></td> <td>Modify Application Definitions</td> </tr> <tr> <td><input type="checkbox"/></td> <td>Delete Application Definitions</td> </tr> <tr> <td><input checked="" type="checkbox"/></td> <td>View Application Definitions</td> </tr> </tbody> </table>	Granted		<input type="checkbox"/>	Create Application Definitions	<input type="checkbox"/>	Modify Application Definitions	<input type="checkbox"/>	Delete Application Definitions	<input checked="" type="checkbox"/>	View Application Definitions
Granted											
<input type="checkbox"/>	Create Application Definitions										
<input type="checkbox"/>	Modify Application Definitions										
<input type="checkbox"/>	Delete Application Definitions										
<input checked="" type="checkbox"/>	View Application Definitions										

Define Role Membership

When you add a [user](#) as a member of a role, the user is granted the access levels defined for the role. If you add a [user group](#) to a role, all users that are members of the user group will automatically be granted the access levels defined for the role.

To define role membership:

1. Select **Access > Roles**.
2. Select the role that you want to add members to from the *Roles* pane. (You can also [create a new role](#).)
3. In the right pane, make sure the **General** tab is selected.
4. Click **Add/Remove**. The *Add/Remove Members* dialog box is displayed.
5. Select the **Users** and/or **User Groups** you want to assign to the role (or deselect users or user groups that you want to remove from the role) and then click **Make Changes**.
6. Click **Save**.

Delete a Role

To delete a role:

1. Select **Access > Roles**.
2. Select the role you want to delete from the *Roles* pane and then click **Remove**.



Manage Devices

A *device* is an external application (such as PolicyCenter), a Blue Coat appliance (PacketShaper), or other generic device (router, bridge, hub, gateway or switch) that you can interact with via IntelligenceCenter. The type of interaction depends on the device type. Some devices, such as PolicyCenter allow [single sign-on](#) functionality only; other devices such as DataCollector allow you to configure how the device works—what type of data it collects, its data retention policies, and which appliances it collects from.

When you log in to IntelligenceCenter for the first time, the [IC network](#) does not contain any device entries. You must create or import the device entries from within IntelligenceCenter before you can begin managing your network devices and collecting data from them. You perform the following device management tasks when managing your network topology:

- [Creating the network groups and sub-groups](#)
- [Adding devices](#)
- [Importing devices from PolicyCenter](#)
- [Viewing and modifying device configurations](#)
- [Removing devices](#)
- [Monitoring devices](#)

In addition, you can create groups of devices — called [views](#) — that represent logical cross-sections of the network. However, you cannot [create network views](#) until you have added the devices to the IC network.


Note: The devices from which you plan to collect data for reporting must have [time synchronization](#) enabled.


IC Network Overview

Before you can start generating IntelligenceCenter (IC) reports, you must configure your IC network. At a minimum, an IC network contains a group of devices and a DataCollector (DC) that collects and reports on the data generated by these devices. Depending on the amount of data you need to collect and the way you want to partition your reporting, you may need to create multiple groups, each with its own DC and unique set of devices. Each DC collects data from and reports on the devices that reside in its own group.

When you first install IC, the network consists of a single group called *Network*. Before you can begin using the IC reporting and single sign-on features, you add the devices you will be collecting data from to the network. To do this, you can either [import devices from a PolicyCenter server](#) or [manually add devices from within IC](#).

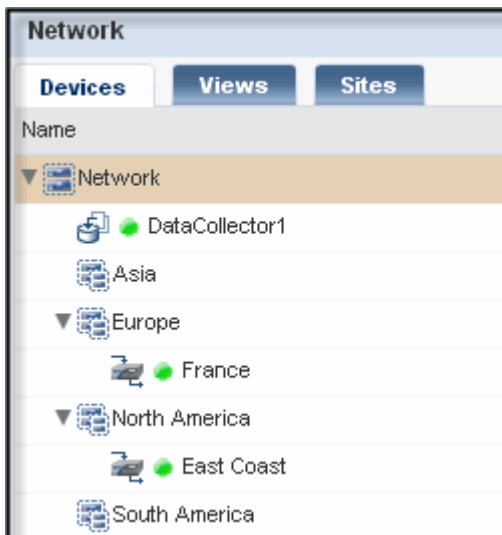
There are two types of objects you can add to the network topology:

- **Network Group** — A top-level organizational unit that defines a reporting domain in IC. Each group must contain a single DC as well as the devices from which this DC will collect and generate reports. A DC in one group cannot collect data from or report on data from devices in a different group. Each group is identified by the  icon.

Within each group that you define, you can further organize your devices into sub-groups, which are represented by the  icon. By organizing small groups of devices into sub-groups, you can optimize the speed at which the network tabs load. You can then report on these sub-groups rather than on the entire network to improve report generation times.



The network groups and sub-groups can be collapsed and expanded to view different sections of the tree — just click the arrow icons next to the folders. The down arrow ▼ icon indicates that the branch is expanded; the right arrow ► icon indicates that the branch is collapsed.



- Device** — An external application or appliance that users can log in to, report on, or access via IC. Devices that can be added to IC include the following BlueCoat products: DataCollector, PacketShaper, and PolicyCenter. In addition, you can add network devices (such as routers or switches). The PolicyCenter and PacketShaper devices that you add to IC can be configured for [single sign-on](#). In addition, the PacketShaper appliances and the network devices such as routers and switches that you add to IC can also be configured as reporting devices. When a device is [configured as a reporting device](#) (called a data source), DataCollector will collect ME and/or FDR data from the device. IC can then display the data in [static reports](#) and [portlets](#).

After you create the IC network, you can [create network views](#), which provide a way to organize devices into multiple logical groupings. You can then use these logical groupings as the basis for reports. For example, you could create a view for each business unit in your organization, allowing you to troubleshoot and monitor performance and applications for each unit independently. Or, you could create views for individual appliances, allowing you to generate reports or portlets based on traffic flowing through that device only.



Create a Group

By default, IC contains a single group called *Network*. You must create a separate group for each DC you plan to deploy.

To create a group:

1. Select **Configure > Network > Devices**.

Note: You can rename the default Network group, but you cannot delete it. If you have a naming scheme for the groups you plan to create, you may want to rename this top-level group to match your scheme. To rename the default group, simply select it in the **Devices** tab, enter a new **Name** and click **Save**.

2. Click **Add** and select **Group** from the pop-up menu. The *Create Group* dialog box is displayed.
3. Enter a **Name** and optionally a **Description** for the group.
4. Specify whether you want to associate a DC with this group now or later:
 - If you want to associate a DC with this group later, select **Do not assign a DataCollector at this time**. You must [add a DC](#) to the group before you can collect data and report on devices in the group.
 - If you have a DC in a different group that you would like to reassign to this new group, select **Reassign a DataCollector**. Note that if you have used this DC to collect data for a different group, you will no longer be able to report on the data you have collected (unless you move the devices as well).
 - If you want to add a new DC to this group, select **Add a new DataCollector** and then [complete the fields required to add the DC](#).
5. Click **Save**. The new group is displayed on the **Devices** tab and is identified with the  icon. Additionally, if you chose to add a DC at the same time that you created the group, the DC is also displayed under the new group object on the **Devices** tab as indicated by the  icon. You must [add a DC](#) to each group in order to collect and report on data for the devices in the group.

Create a Sub-Group


If you have many devices, you can optimize the speed at which the network tabs load by organizing small groups of devices into sub-groups. You can then report on these sub-groups rather than on the entire network to improve report generation times.

To create a sub-group:

Click **Add** and select **Sub-group** from the pop-up menu. The *Create Sub-group* dialog box is displayed..


To create a sub-group:

1. Select **Configure > Network**.
2. If it's not already selected, select the **Devices** tab.
3. Select the group or sub-group under which you want to create the new sub-group. The network groups can be collapsed and expanded; just click the arrow icons next to each group or sub-group. The down arrow ▼ icon indicates that the group is expanded; the right arrow ► icon indicates that the group is collapsed.



4. Click **Add** and select **Sub-group** from the pop-up menu. The *Create Sub-group* dialog box appears.

The screenshot shows a dialog box titled 'Create Sub-group' with a help icon in the top right. It contains two input fields: 'Name' with a red asterisk indicating it is required, and 'Description'. At the bottom, there are two buttons: 'Save' and 'Cancel'.

5. In the *Create Sub-group* window, enter the **Name** and optionally a **Description** for the network branch and click **Save**. The new sub-group is displayed in the topology and is identified with the  icon.

Add a Device


In order to interact with a device from within IntelligenceCenter (IC), the device must have an entry in the [network topology](#). The network *topology* represents the physical view of your network. When you add a device to IC, you must add it to the topology based on its physical location in your network. Note that you should [create the network branches](#) that represent the geographical locations on your network before you begin adding devices.

You can either [import the devices from a PolicyCenter installation](#) or you can manually add the devices using IC.


Note: The number of devices that you can add to IC is controlled by your [product license](#).

To manually add a device to IC:

1. Select **Configure > Network > Devices**.
2. Select the [network group or sub-group](#) where the device you are adding is located and click **Add**.



3. Select the device type from the pop-up menu.



4. Configure the device. The configuration information you supply depends on the type of device you are adding:
 - [DataCollector](#)
 - iShaper (no longer supported)
 - [Network Device](#)
 - [PacketShaper](#)
 - [PolicyCenter](#)
5. When you are done configuring the device, click **Save**. The new device is displayed in the topology.

Configure a Device

Configure a PacketShaper Device Entry


You must complete the following fields when [adding](#) or [modifying](#) a PacketShaper device entry:

Note: IC supports PacketShaper appliances running PacketWise version 7.3.1 or higher.

Field	Description
Device Name	A descriptive name for the PacketShaper appliance.
Description	An optional description of the device.
Serial Number *	Serial number of the PacketShaper appliance.
Host *	IP address or DNS hostname of the PacketShaper appliance.
Secure Login	Indicates whether communication with this PacketShaper appliance —both for status updates to IntelligenceCenter and ME data collection from DataCollector—will happen over a secure SSL connection.
Port	The port number on which IntelligenceCenter and DataCollector will communicate with this PacketShaper appliance.
Touch Password	The password required by the PacketShaper appliance before read/write access is granted. Note that if you change this password on the PacketShaper appliance after adding it to IC, you will have to edit this field with the new password in order for IC to continue to interact with the appliance.
Location	An optional description of the device location.

Note that required fields are marked with a red asterisk (*). If you do not complete a required field or if you enter an incorrect value, the field text box will be outlined in red. After you complete all required fields, click **Save**.

After you save the PacketShaper device entry, the following additional information is displayed in the **General** tab in the right pane:

Field	Description
Device Type	The type of PacketShaper appliance.
Software Version	The version of PacketWise that is running on the PacketShaper.
Time Zone	The time zone that is configured on the PacketShaper.
Link Size	The size of the PacketShaper appliance's inbound and outbound links in bytes.
Last Contact	The date and time that IntelligenceCenter last successfully pinged the PacketShaper.
Last Update	The last time the information displayed on this screen was last updated. To update the information, you must manually press the update  button. If you have not clicked update, this field will display <i>unavailable</i> .
Connection Status	<p>Indicates whether IntelligenceCenter is able to ping the PacketShaper appliance. There are three possible states:</p> <ul style="list-style-type: none"> • Connected — IntelligenceCenter has full connectivity to the PacketShaper. • Reachable Not Connected — IntelligenceCenter can ping the PacketShaper appliance, but is unable to communicate with the device. This may indicate that some of the field values, such as serial number or touch password, were entered incorrectly. • Unreachable — IntelligenceCenter cannot ping the PacketShaper appliance.
Services	
Shaping	Indicates whether shaping is enabled on the PacketShaper appliance.
Traffic Discovery	Indicates whether traffic discovery is enabled on the PacketShaper appliance.
Acceleration	Indicates whether acceleration is enabled on the PacketShaper appliance.
Compression	Indicates whether compression is enabled on the

	PacketShaper appliance.
Single Sign-On Credential Information	
User Name	The user name required to log in to the PacketShaper appliance, if any. Each IntelligenceCenter user must enter his own user name in order to enable single sign-on to this device. If there is no user name, it is okay to leave this field blank.
Password	The password required to log in to the PacketShaper appliance. Each IntelligenceCenter user must enter his own password in order to enable single sign-on to this device.

General

Configurations

Device Name

* France

Description

Device Type

PacketShaper 2500

Serial Number

* 025-10002432

Host

10.9.50.91

[Sign On](#)

Port

80

Secure Login

☐

Touch Password

Location

Software Version

8.4.1b7X

Time Zone

America/Los_Angeles

Link Size

40 Mbps inbound 60 Mbps outbound

Last Contact

Aug 9, 2010 12:28 PM

Last Update

Oct 21, 2009 11:46 AM

Connection Status

Connected

Services

Shaping

Off

Acceleration

Unavailable

Traffic Discovery

On

Compression

Off

Single Sign-On Credential Information

User Name

Password

Configure PolicyCenter

You must complete the following fields when [adding](#) or [modifying](#) a PolicyCenter device entry:

Field	Description
Device Name	A descriptive name for the PolicyCenter application
Description	An optional description of the application
Host *	IP address of the PolicyCenter Directory Server
Port	Port number to be used for web access (default = 80)
Secure Login	Select this checkbox if the specified port is secure (such as port 443)
Administrator ID	The user name for the PolicyCenter administrative user
Administrator Password	The password associated with the PolicyCenter administrator ID
Location	An optional description of the PolicyCenter server location
<i>In order to enable single sign-on, each user must specify the following login information after PolicyCenter is successfully added to the network topology:</i>	
User Name	User name for the PolicyCenter user
Password	Password for the specified user account

Note that required fields are marked with a red asterisk (*). If you do not complete a required field or if you enter an incorrect value, the field text box will be outlined in red. After you complete all required fields, click **Save**.

After IntelligenceCenter successfully connects with PolicyCenter, it retrieves the device type, serial number and software version as shown below. In addition, it provides fields for users to [configure single sign-on](#) as well as a *Import Shaper Inventory* section that you can use to [import the PacketShaper appliances](#) managed by the PolicyCenter server into your IntelligenceCenter network topology. The **Last Update** field indicates the date and time the information displayed on this screen was last updated. To update the information, you must manually click the **Sync** button. The **Last Contact** field indicates the date and time that IntelligenceCenter last successfully pinged PolicyCenter.

10.9.59.73

?

General Information

Device Name

10.9.59.73

Description

Device Type

PolicyCenter

Serial Number

901-20000100

Software Version

8.5.3b13

Host

10.9.59.73

[Sign On](#)

Port

80

Secure Login

☐

Location

Last Contact

Dec 8, 2010 12:58 PM

Last Update

Dec 8, 2010 12:58 PM

Administrator ID

Administrator Password

Single Sign-On Credential Information

User Name

Password

Import Shaper Inventory

Import Shapers

Save

Cancel

Sync

Configure DataCollector

You must complete the following fields when [adding](#) or [modifying](#) a DataCollector device entry:

Field	Description
Device Name	A descriptive name for DataCollector
Description	An optional description of the device.
Host	IP address or DNS hostname of the server on which DC is installed. If IC and DC are installed on the same system, you can enter localhost in this field.
Port	Port number to be used for communication between DataCollector and IntelligenceCenter. The default port (8543) cannot be modified.
Secure Login	Indicates that the communication between DataCollector and IntelligenceCenter occurs over a secure port.
Touch Password	The DataCollector password created when you installed DataCollector.
Location	An optional description of the server location

Note that required fields are marked with a red asterisk (*). If you do not complete a required field or if you enter an incorrect value, the field text box will be outlined in red. After you complete all required fields, click **Save** to save your changes.

After IntelligenceCenter successfully connects with DataCollector, it retrieves the device type, serial number, and software version. In addition, four tabs are displayed in the right-hand pane: **General**, **Status**, **Configuration**, and **Data Sources**. You use these tabs to modify the data collection settings and [monitor the overall health of DC](#). In order to enable data collection, you must also [define data sources](#).

DataCollector1

General

Status

Configuration

Data Sources

Name

* DataCollector1

Description

Device Type

DataCollector

Port

8543

Secure Login

☒

Host

localhost

Location

Serial Number

950-14233383

Software Version

collector3.0.0b308086

Last Contact

Aug 9, 2010 12:13 PM

Last Update

Oct 21, 2009 11:21 PM

Connection Status

Connected

Discover Traffic Classes

Configure a Network Device

You can add network devices such as routers, switches, or hubs as IntelligenceCenter devices. You can then configure these devices to emit flow records to DC and use the data you collect to generate reports.

You must complete the following fields when [adding](#) or [modifying](#) a network device entry:

Field	Description
Device Name	A descriptive name for the device
Description	An optional description of the device
Device Type	Specifies the type of network device. Select one of the following values from the drop-down list: Router , Gateway , Bridge , Switch , Hub , or Other .
Host *	IP address or DNS hostname of the device
Port	Port number to be used for web access (default = 80)
Secure Login	Select this checkbox if the specified port is secure (such as port 443)
Location	An optional description of the device location

Note that required fields are marked with a red asterisk (*). If you do not complete a required field or if you enter an incorrect value, the field text box will be outlined in red. After you complete all required fields, click **Save**.

Import Devices from PolicyCenter

As an alternative to adding PacketShaper appliances to your IntelligenceCenter [IC network](#) manually, you can import device entries from a PolicyCenter server. Note that you must [add PolicyCenter to the IC network](#) and define valid PolicyCenter administrative credentials in its IC device entry before you can import appliances from it.

Note: All devices that are imported from PolicyCenter are added to the top-level group; you cannot import devices into a different sub-group. However, you can move the devices to a different sub-group after they are imported.

To import devices from PolicyCenter:

1. Select **Configure > Network > Devices**.
2. Select the network group where the PolicyCenter server is located:
 - If you have already added the PolicyCenter server to IC, select its entry on the **Devices** tab. The PolicyCenter server details are displayed in the right pane.
 - If you have not yet added the PolicyCenter server to IC, click **Add** and select **PolicyCenter** from the pop-up menu. Enter the required PolicyCenter configuration information and then click **Save**. Details about the PolicyCenter server you just imported are displayed in the right pane.

PC_132

General Information

Device Name: PC_132

Description:

Device Type: PolicyCenter

Serial Number: 901-20000100

Software Version: 8.3.1g1

Host: 10.9.88.132 [Sign On](#)

Port: 80 Secure Login ☐

Location:

Administrator ID *: admin

Administrator Password: *****

Single Sign-On Credential Information

User Name:

Password:

Import Shaper Inventory

Import Shapers

3. In the *Import Shaper Inventory* section of the right pane, click **Import Shapers**. IntelligenceCenter begins importing the devices from PolicyCenter; this process may take several minutes depending on the number of devices being imported. When the import completes, a message displays telling you that you must refresh the screen in order to see the imported devices. After you refresh the screen (which requires you to log out and log back in), the new devices are displayed just below the top-level Group object (called *Network* by default); each appliance is uniquely identified by its serial number.
4. If you want to move the newly imported devices to a different sub-group, you can drag the device icon to the new location on the **Devices** tab.
5. To change the name that is displayed for an imported device from a serial number to a more descriptive name, click on the device entry on the **Devices** tab, enter a new **Device Name** in the **General** tab in the right pane and then click **Save**.

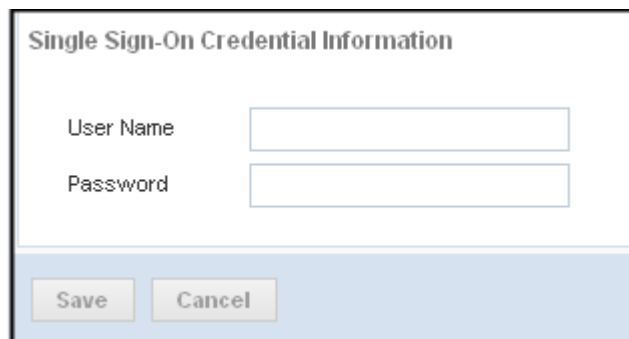
Configure Single Sign-On

After a Blue Coat device has been added to the IntelligenceCenter topology, IntelligenceCenter users can enter their own sign-on credentials for the devices they use. Users can then use the [IC network topology](#) as their dashboard for launching applications. With a single click, users can log in to PolicyCenter or any PacketShaper appliances — without entering a user name and password.

Note: Users can only use the single sign-on feature to access appliances or applications on which they have a local user name and password configured. Users who access the appliance or software via a RADIUS or TACACS+ account cannot use the single sign-on feature.

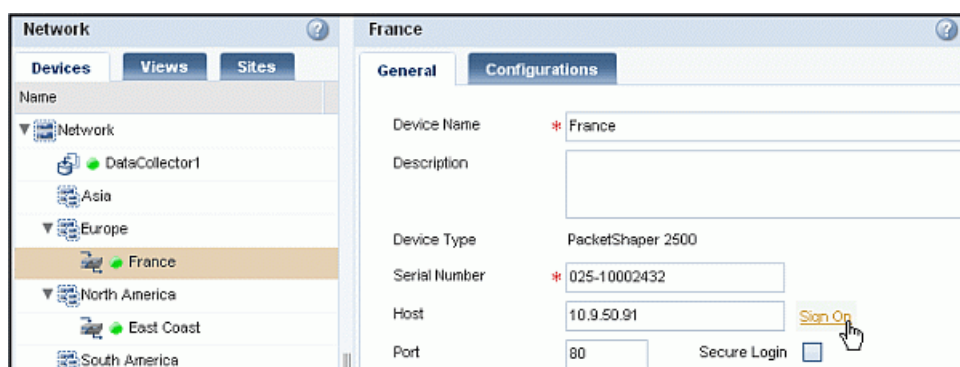
To enable single sign-on, each user must do the following:

1. [Log in to IntelligenceCenter](#).
2. Select **Configure > Network > Devices**.
3. Navigate to the device on which to configure single sign-on. The network groups can be collapsed and expanded; just click the arrow icons next to each branch. The down arrow ▼ icon indicates that the group is expanded; the right arrow ► icon indicates that the group is collapsed.
4. Select the device. The device information is displayed in the right-hand pane. Locate the *Single Sign-On Credential Information* section and enter the **User Name** and **Password** as necessary.










The dialog box titled "Single Sign-On Credential Information" contains two input fields: "User Name" and "Password". Below these fields are two buttons: "Save" and "Cancel".

5. Click **Save**.
6. Each user can now sign on to the device from the **Devices** tab in one of the following ways:
 - In [Details view](#), select the device you want to sign on to in the network topology and then click the **Sign On** link on the **General** tab:



- In [Summary view](#), click the IP address or hostname of the device in the **Host** column to sign on to it.

Devices										
Views		Sites								
Name	Host	Type	Serial Number	Software Version	Acceleration	Compression	Shaping	Discovery	Last Update	
▼ Network										
 DataCollector-Bulk50218	127.0.0.1	DataCollector	950-01348661	collector3.1.1b50218						Sep 15, 2010 11:37 AM
 10.78.53.18-S	10.78.53.18	PacketShaper 2500	025-10000250-S	8.5.1g1	Off	Off	Off	On		Sep 15, 2010 9:09 AM
 Lab-PS10k	10.78.53.182	PacketShaper 10000	011-10005014	8.5.4g1	Off	Off	On	Off		Sep 12, 2010 10:23 PM
 Lab-PS10K-2	10.78.53.181	PacketShaper 10000	011-10001009	8.4.6g1	Off	Off	Off	On		Sep 12, 2010 10:42 PM
 PS-KL	10.105.13.166	PacketShaper	175-10000962			Off	Off			
 PS10.78.52.47	10.78.52.47	PacketShaper 10000/ISP	011-10005485	8.5.4g1		Off	On	Off		Sep 15, 2010 11:38 AM
 Sunnyvale	10.78.53.135	PacketShaper 1550	015-10900526	8.4.4g1	On	Off	Off	Off		Sep 12, 2010 10:34 PM

Modify a Device Configuration

To view and/or modify a device configuration:

- Select **Configure > Network > Devices**.
- Locate the device in the network. The network groups can be collapsed and expanded to help you find the device you're looking for. Just click the arrow icons next to each branch. The down arrow ▼ icon indicates that the group is expanded; the right arrow ► icon indicates that the group is collapsed.
- If you want to move the device to a different group or sub-group in the network, you can drag the device icon to the new location on the **Devices** tab.
- If you are in Summary view, click the **Details View** link to switch views.

Network

Devices Views Sites

Name	Host	Type	Serial Number	Software Version	Acceleration	Compression	Shaping	Discovery	Last Update
Network									
DataCollector-Build50218	127.0.0.1	DataCollector	950-01346661	collector3.1.1b50218					Sep 15, 2010 11:37 AM
10.78.53.18-S	10.78.53.18	PacketShaper 2500	025-10000250-S	8.5.1g1	Off	Off	Off	On	Sep 15, 2010 9:09 AM
Lab-PS10k	10.78.53.182	PacketShaper 10000	011-10005014	8.5.4g1	Off	Off	On	Off	Sep 12, 2010 10:23 PM
Lab-PS10K-2	10.78.53.181	PacketShaper 10000	011-10001009	8.4.6g1	Off	Off	Off	On	Sep 12, 2010 10:42 PM
PS-KL	10.105.13.166	PacketShaper	175-10000962			Off	Off		
PS10.78.52.47	10.78.52.47	PacketShaper 10000/ISP	011-10005485	8.5.4g1		Off	On	Off	Sep 15, 2010 11:38 AM
Sunnyvale	10.78.53.135	PacketShaper 1550	015-10900526	8.4.4g1	On	Off	Off	Off	Sep 12, 2010 10:34 PM

88

100

Add Remove

Show Details Sync

- Select the device you want to view. When you select the device, the corresponding configuration information is displayed in the right-hand pane.

The screenshot shows a software interface for managing a network. On the left, a tree view under 'Network' shows a hierarchy: Network > Europe > France (selected). Other nodes include DataCollector1, Asia, North America, East Coast, and South America. The main panel on the right is titled 'France' and has two tabs: 'General' (active) and 'Configurations'. The 'General' tab contains the following fields and values:

- Device Name: * France
- Description: (empty text box)
- Device Type: PacketShaper 2500
- Serial Number: * 025-10002432
- Host: 10.9.50.91 (with a 'Sign On' link)
- Port: 80 (with a 'Secure Login' checkbox)
- Touch Password: (empty text box)
- Location: (empty text box)
- Software Version: 8.4.1b7X
- Time Zone: America/Los_Angeles
- Link Size: 40 Mbps inbound 60 Mbps outbound
- Last Contact: Aug 10, 2010 1:23 PM
- Last Update: Oct 21, 2009 11:46 AM
- Connection Status: Connected

Below these fields is a 'Services' section with a table:

Service	Status	Feature	Status
Shaping	Off	Acceleration	Unavailable
Traffic Discovery	On	Compression	Off

At the bottom of the 'General' tab is a 'Single Sign-On Credential Information' section with fields for 'User Name' and 'Password'.

At the bottom of the interface are buttons: '+ Add', '- Remove', 'Summary View', 'Save', 'Cancel', and a refresh icon.

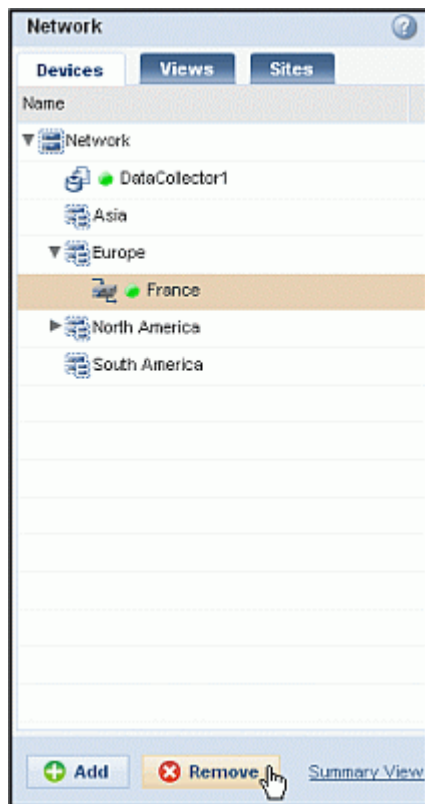
6. Modify the device configuration as desired. The configuration information that is displayed depends on the type of device you selected:
 - [PacketShaper](#)
 - iShaper (no longer supported)
 - [DataCollector](#)
 - [PolicyCenter](#)
 - [Network Device](#)
7. When you are done configuring the device, click **Save**.

Remove a Device

If you no longer want to interact with a device from within IntelligenceCenter, you can remove it from the network topology. If you are collecting data from the device, you must [remove it as a data source](#) before you can delete it.

To remove a device from the IC network:

1. Select **Configure > Network > Devices**.
2. Locate the device you want to remove. The network groups and sub-groups on the **Devices** tab can be collapsed and expanded to help you find the device you're looking for. Just click the arrow icons next to each branch. The down arrow ▼ icon indicates that the branch is expanded; the right arrow ► icon indicates that the branch is collapsed.
3. Select the device you want to remove and then click **Remove**.



4. When prompted to confirm the device removal, click **Yes**.

Manage PacketShaper Configurations

PacketShaper settings are stored in a configuration file named `config.ldi`. This file contains the traffic tree configuration (including all classes, class IDs, partitions, policies, host lists, and events), as well as all sharable configuration settings such as packet shaping, traffic discovery, passwords, SNMP, email, SNTP, compression, and Syslog. The configuration file should be backed up on a regular basis, as it can be used to restore a configuration if needed. Although you can back up and restore configurations using CLI commands on the PacketShaper itself, IntelligenceCenter offers an easier way to do this via the PacketShaper device's **Configurations** tab.

- [Back Up a PacketShaper Configuration](#)
- [Restore a PacketShaper Configuration](#)
- [View a PacketShaper Configuration](#)
- [Delete a PacketShaper Configuration](#)
- [Compare PacketShaper Configurations](#)

Back Up a PacketShaper Configuration

There are several reasons why you would want to back up your PacketShaper configuration:

- Your configuration may not load properly after upgrading to a new version of PacketWise software.
- You could upgrade to a new release and then realize that you want to revert to a previous version.
- To recover critical configuration files in the unlikely event of file corruption.

To back up a PacketShaper configuration:

1. Select **Configure > Network > Devices**.
2. Select the [network group or sub-group](#) where the PacketShaper device is located and select the PacketShaper device name. The PacketShaper details are displayed in the right pane.
3. Select the **Configurations** tab in the right pane.
4. Select the configuration you want to back up (for example, Current).
5. Click **Backup**. The *Backup Configuration* dialog box displays.
6. In the **Name** field, enter a name for the configuration. The name can be up to 32 characters long; spaces and special characters are allowed.
7. (Optional) In the **Description** field, enter a brief description of the configuration. The description can be up to 64 characters long; spaces and special characters are allowed.
8. Click **Backup**. The **Configurations** tab now displays the configuration name and description, along with the user who backed up the configuration and the current date and time. Note that the configuration is stored in the IntelligenceCenter database, not on the PacketShaper itself.

Restore a PacketShaper Configuration

In the event that your current PacketWise configuration becomes corrupt or you want to roll back to a previously saved configuration, you can restore any PacketShaper configuration that you have backed up in IC.

To restore a PacketShaper configuration:

1. Select **Configure > Network > Devices**.
2. Select the [network group or sub-group](#) where the PacketShaper device is located and select the PacketShaper device name. The PacketShaper details are displayed in the right pane.
3. Select the **Configurations** tab in the right pane.
4. In the list of backed up configurations, select the name of the configuration you want to restore.

Note: You cannot select **Current** because that is the name of the configuration currently in use.

5. Click **Restore**. After a moment you will get a confirmation message that the configuration was restored.
6. Click **OK**.

Notes:

- The IC **Configurations** tab does not offer a way to restore a configuration from another PacketShaper; if you need to do this, use the PacketShaper CLI. For details on this procedure, refer to [Back Up a Configuration](#) in PacketGuide.
- Because PolicyCenter has its own way of managing configurations, you cannot restore configurations to a PacketShaper that is in shared mode (that is, prescribed to PolicyCenter). The PacketShaper must be in local mode.

View a PacketShaper Configuration

If a configuration's name, description, and date don't provide sufficient clues as to what is contained in the configuration, you can view its contents. When you view a PacketShaper configuration, the information is presented in a directory schema format showing configuration objects and their attributes. Many of the objects are self-explanatory — for example, `iqosShapingOn: off` indicates that shaping is turned off in the configuration and `iqosFlowDetailRecordsCollectorIP: 10.9.50.92` specifies the IP address of an FDR collector. Don't be concerned with understanding all of the directory schema as some of the language is a bit cryptic; however, you can usually pick out key settings based on the descriptive object names.

To view a configuration:

1. Select **Configure > Network > Devices**.
2. Select the [network group or sub-group](#) where the PacketShaper device is located and select the PacketShaper device name. The PacketShaper details are displayed in the right pane.
3. Select the **Configurations** tab in the right pane.
4. In the list of backed up configurations, select the name of the configuration you want to view.

5. Click **View**. The objects and attributes of the selected configuration appear in a scrollable list.

Tip: If you would like to see which settings are different between two configurations, you can [compare](#) them.

Delete a PacketShaper Configuration

Once a PacketShaper backup configuration has become outdated or is no longer needed, you can delete it. Note that you cannot undelete a configuration or undo the Delete command.

To remove a PacketShaper backup configuration:

1. Select **Configure > Network > Devices**.
2. Select the [network group or sub-group](#) where the PacketShaper device is located and select the PacketShaper device name. The PacketShaper details are displayed in the right pane.
3. Select the **Configurations** tab in the right pane.
4. In the list of backed up configurations, select the name of the configuration you want to delete.

Note: You cannot select **Current** because that is the name of the configuration currently in use.

The configuration is deleted.

Compare PacketShaper Configurations

If you would like to see which settings are different between two configurations, you can compare them.

To compare two configurations:

1. Select **Configure > Network > Devices**.
2. Select the [network group or sub-group](#) where the PacketShaper device is located and select the PacketShaper device name. The PacketShaper details are displayed in the right pane.
3. Select the **Configurations** tab in the right pane.
4. In the list of backed up configurations, select the name of the first configuration you want to compare and then hold down Ctrl as you click the name of the second configuration.
5. Click **Diff**. The two configurations are displayed side by side; the objects and attributes that differ between the configurations are highlighted in different colors.

Monitor Devices

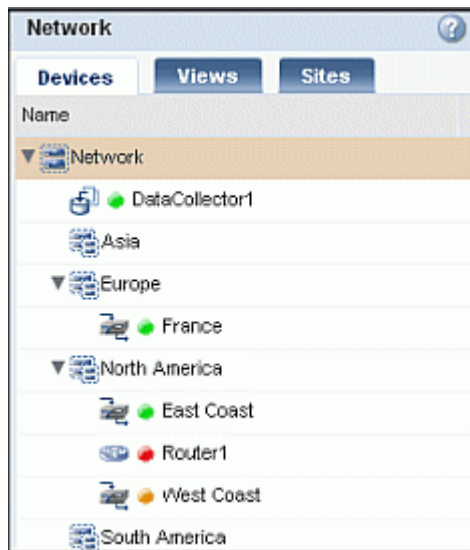
A *device* is an external application or appliance that you can interact with via IntelligenceCenter (IC). After you add devices to IC, you can monitor them as follows:

- You can [view summary and detailed information](#) about all devices that have been added to the IC network.
- You can [monitor DataCollector health](#).
- You can [view the class tree](#) of any PacketShaper appliance that you have added as an ME data source.

View Device Information

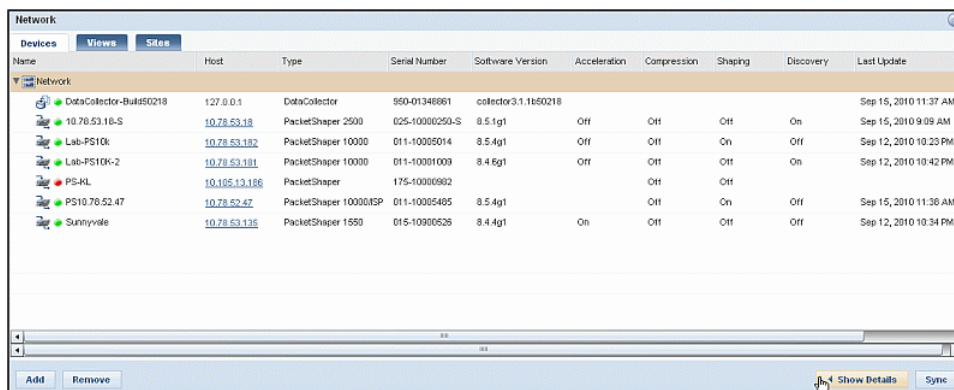
To monitor the devices you have added to IC:

1. Select **Configure > Network > Devices**. The network groups and sub-groups you have created are displayed; you can collapse and expand the groups to help you find the devices you're looking for. Just click the arrow icons next to each group. The down arrow ▼ icon indicates that the group is expanded; the right arrow ► icon indicates that the group is collapsed.
2. Check the health of your devices by looking at the colored icon next to the device IP address or hostname in the **Name** field on the **Devices** tab: a green ● icon means that IC can ping the device and log in to it; a yellow ● icon means that IC can ping the device, but cannot log in to it; a red ● icon indicates that IC can neither ping nor log in to the device.



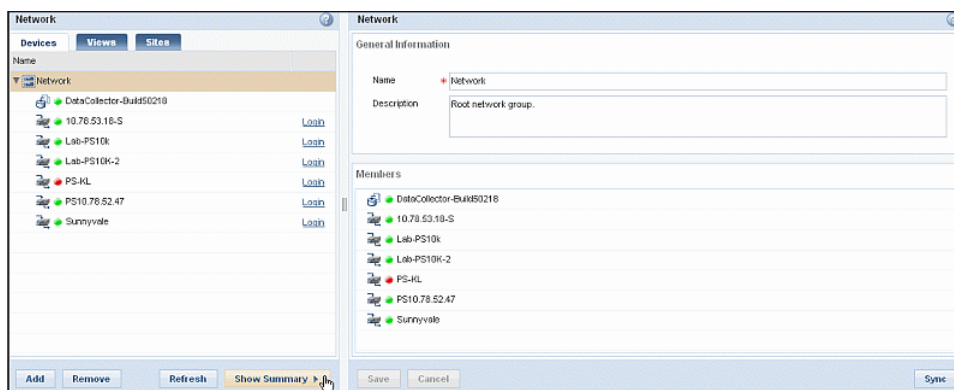
3. View information about the devices by toggling between Summary view and Details view as follows:
 - To get a quick picture of the status of all of your devices, click the **Show Summary** link at the bottom of the *Network* window. In Summary view you can quickly see identifying information for the device (device type, IP address or hostname, serial number, and software version), services that are running on

each device (acceleration, compression, and/or shaping), as well as information about whether class discovery is on and the date and time that IC was last in contact with the device.



Name	Host	Type	Serial Number	Software Version	Acceleration	Compression	Shaping	Discovery	Last Update
DataCollector-Build50218	127.0.0.1	DataCollector	950-01348661	collector3.1.1b50218					Sep 15, 2010 11:37 AM
10.78.53.18-S	10.78.53.18	PacketShaper 2500	025-10000250-S	8.5.1g1	Off	Off	Off	On	Sep 15, 2010 9:09 AM
Lab-PS10k	10.78.53.182	PacketShaper 10000	011-10005014	8.5.4g1	Off	Off	On	Off	Sep 12, 2010 10:23 PM
Lab-PS10K-2	10.78.53.181	PacketShaper 10000	011-10001009	8.4.6g1	Off	Off	Off	On	Sep 12, 2010 10:42 PM
PS-KL	10.105.13.186	PacketShaper	175-10000602		Off	Off	Off		
PS10.78.52.47	10.78.52.47	PacketShaper 10000ISP	011-10005405	8.5.4g1		Off	On	Off	Sep 15, 2010 11:36 AM
Sunnyvale	10.78.53.135	PacketShaper 1550	015-10900528	8.4.4g1	On	Off	Off	Off	Sep 12, 2010 10:34 PM

- To view more detailed information about a specific device, click the **Show Details** link at the bottom of the *Network* window and then select the device you want to monitor from the network topology. When you select a device, details about the device are displayed in the right-hand pane. The details that are displayed depends on the type of device you selected. For example, if you select a [PacketShaper](#), two tabs (**General** and **Configurations**) are displayed. If you select [DataCollector](#), four tabs (**General**, **Status**, **Configuration**, and **Data Sources**) are displayed.



The screenshot shows the 'Network' window with the 'Views' tab selected. On the left, a list of devices is shown, including 'DataCollector-Build50218', '10.78.53.18-S', 'Lab-PS10k', 'Lab-PS10K-2', 'PS-KL', 'PS10.78.52.47', and 'Sunnyvale'. On the right, the 'General Information' tab is active for the selected device. It displays the device's name, description, and a list of members.

General Information

Name: DataCollector-Build50218

Description: Root network group.

Members

- DataCollector-Build50218
- 10.78.53.18-S
- Lab-PS10k
- Lab-PS10K-2
- PS-KL
- PS10.78.52.47
- Sunnyvale

- To refresh the information that is displayed on either the Summary view or the Details view, click **Sync** in the lower right of the screen. You must manually refresh the data in order to see up-to-date information; the **Last Update** field indicates the time at which data was last fetched from the device.

Monitor DataCollector Health


To monitor DataCollector:

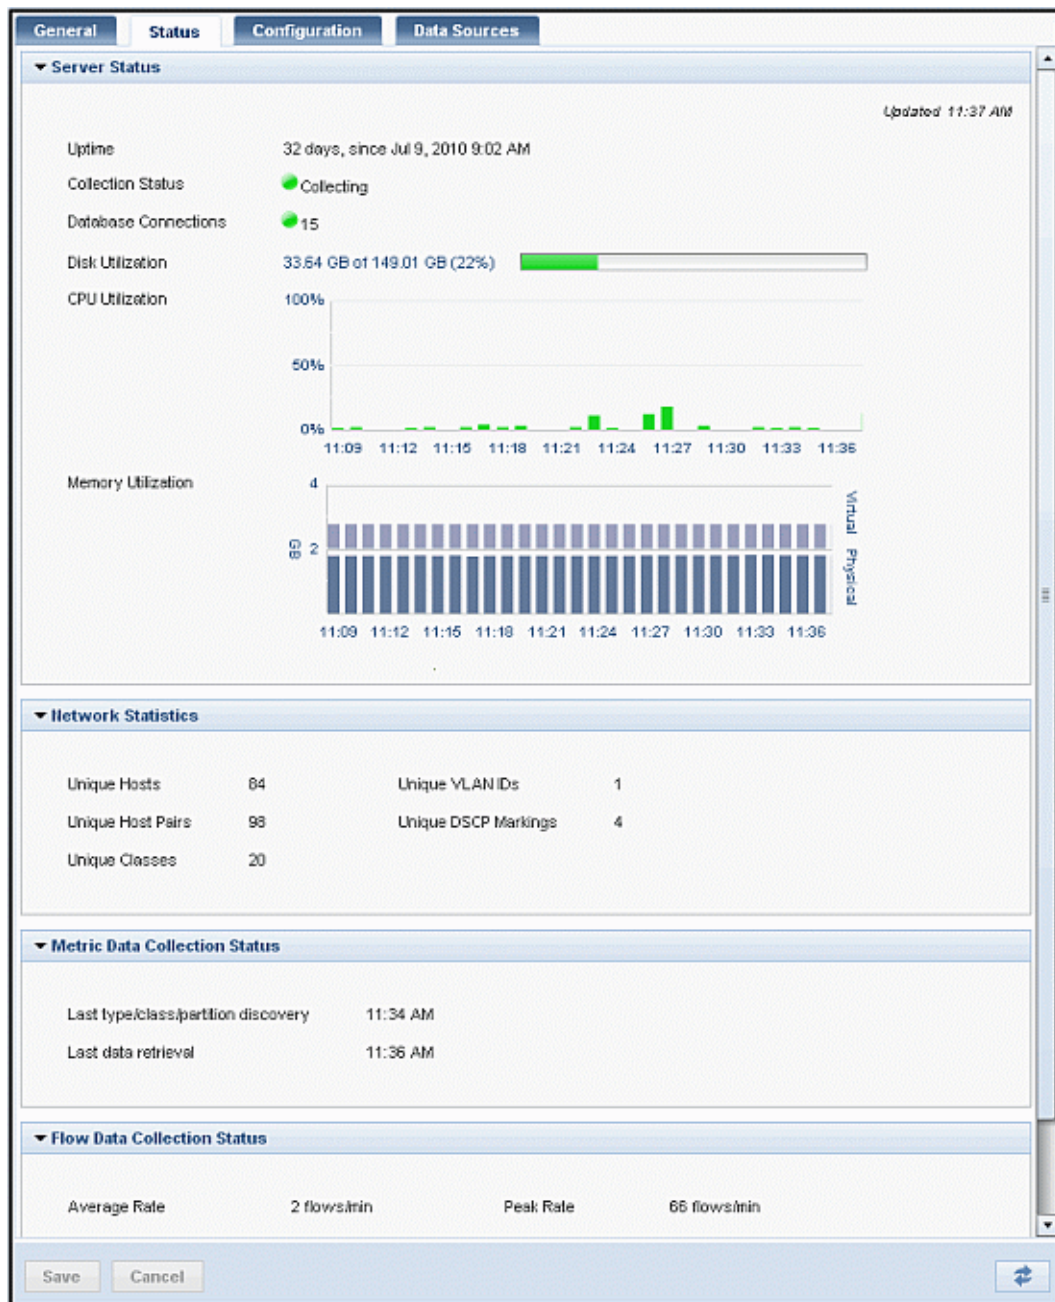
1. Select **Configure > Network > Devices**.
2. Select the DataCollector object on the **Devices** tab. The DataCollector is indicated by the  icon.

Note: To expand or collapse a network branch, click the arrow icons next to each branch. The down arrow ▼ icon indicates that the branch is expanded; the right arrow ► icon indicates that the branch is collapsed.

When you select the object, the DataCollector **Status** tab is displayed.





This screen displays health information about the server on which DataCollector is installed, status information about metric and flow data collection, and statistical information about the data that has been collected. In addition, if there are any collection errors, they are displayed at the bottom of the screen.


The status information on this screen is dynamically updated (every 60 seconds by default unless you've changed the status interval in [your IC profile](#)). You can also click the update  button in the lower right of the screen to force IntelligenceCenter to update the status now.



This tab displays the following information:

Field	Description
Server Status	
Uptime	The number of days since and the date and time that DataCollector last started.

Collection Status	<p>Indicates whether data collection has been enabled. Possible values are:</p> <ul style="list-style-type: none">  Collecting — DataCollector has started and is ready to collect data. Note that this does not indicate that DataCollector is actually collecting any data. See the <i>Metric Collection Status</i> and <i>Flow Data Collection Status</i> sections of this screen to determine whether data is being collected.  Stopped — DataCollector has stopped and is not collecting data. Possible reasons include license expiration, mismatched versions of IC and DC, or disk utilization reaching full capacity.  Migrating — DataCollector is migrating the database from a previous software release.  Initializing — DataCollector is starting up. You should only see this status for the first few minutes after a restart of the system.
Database Connections	<p>The number of open connections to DataCollector's Postgres database. Connections get opened whenever DataCollector inserts data or whenever a user performs report or portlet operations from IntelligenceCenter.</p>
Disk Utilization	<p>The amount of disk space used by the DataCollector database. This field shows both the amount and percentage of disk space used and the total available disk space. It also shows a status bar that indicates whether you are in danger of running out of disk space. When the database reaches 95% of its disk space capacity, DC will stop collecting data and reporting will be interrupted. If you notice that disk utilization is nearing the 95% mark, you should take action to free up disk space on the drive before collection stops.</p>
CPU Utilization	<p>Shows CPU utilization over the last 10 minutes on the server where DataCollector is installed.</p>

Memory Utilization	Shows virtual and physical memory utilization over the last 10 minutes on the server where DataCollector is installed.
Network Statistics <i>This section shows statistics about the flow data that DataCollector has collected. If you are not collecting FDR, these values will be 0.</i>	
Unique Hosts	The number of unique host entries in the DataCollector database.
Unique Host Pairs	The number of unique host pair entries in the DataCollector database.
Unique VLAN IDs	The number of unique VLAN IDs in the DataCollector database.
Unique Classes	The number of unique PacketShaper traffic classes in the DataCollector database.
Unique DSCP Markings	The number of unique Differentiated Services Code Point (DSCP) values in the DataCollector database.
Metric Data Collection Status <i>This section shows information about ME data collection and indicates if there are errors with your ME data sources.</i>	
Last link/class/partition discovery	Date and time that DataCollector last discovered traffic class and partition data from its data sources.
Last data retrieval	Date and time that DataCollector last collected ME data from one of its ME data sources.
x of y Configured data sources have errors	<p>This section is only displayed if one or more of your ME data sources has errors, where x indicates the number of configured ME data sources with errors and y indicates the total number of configured ME data sources.</p> <p>Each error lists the IP address of the Shaper that had the error and provides a brief description. To view ME data collection statistics for the device, click the right-arrow  icon next to error message entry. You can also click the IP address to sign in to the appliance and conduct further troubleshooting.</p>

Flow Data Collection Status <i>This section shows information about FDR data collection and indicates if there are errors with your FDR data sources.</i>	
Average Rate	The average number of flows per minute received since DataCollector last started.
Peak Rate	The peak number of flows per minute received since DataCollector last started.
Invalid Records	The number of invalid FDR records that DataCollector received. Invalid records include records from configured data sources that are in a format not supported by DataCollector, such as intermediate FDR (iFDR).
Rejected Records	The number of FDR records that DataCollector rejected. DataCollector only accepts FDR records from its configured data sources .
Total Records	The total number of flow records collected.
Total Duplicate Records	Of the total flow records collected, the number that DC has determined to be duplicate. Duplicate records may be collected if the same flow travels through multiple PacketShaper appliances.
Flow Data Collection Errors	This section is only displayed if DC has encountered FDR errors. Each FDR error message shows the data and time of the FDR, the type of FDR (Packeteer-2 or NetFlow-5), the IP address of the device that sent the FDR, and a description of the error (for example, <i>Device not configured</i>).

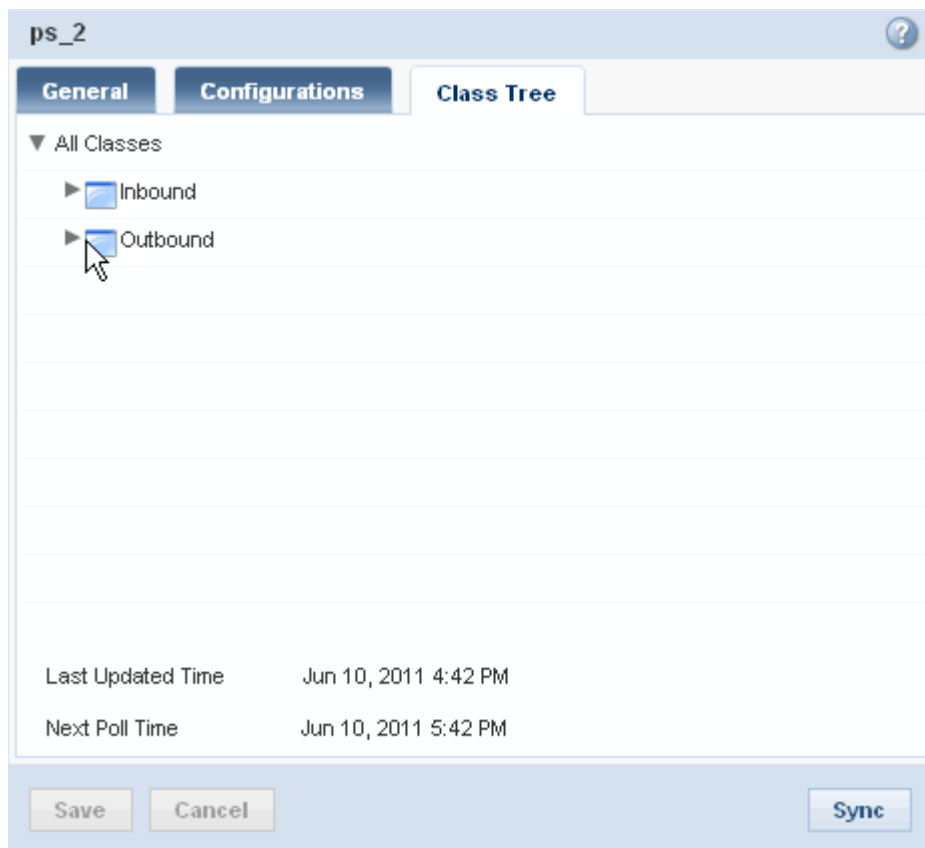
View the PacketShaper Class Tree

For each of the PacketShaper appliances that you have added as ME data sources, you can view the class tree as it appears on the PacketShaper. If the PacketShaper is a ME data source, DataCollector will discover the traffic classes during its regular collection intervals (or on demand). IC will then display the discovered class data in a hierarchical tree structure on the **Class Tree** tab.

To view the class tree:

1. Select **Configure > Network**.
2. On the **Devices** tab, select the PacketShaper appliance for which you want to display the class tree.
3. Select the **Class Tree** tab. The class tree displays along with the time the tree was last updated and the time the tree will next be updated (**Next Poll Time**).

Note: The Last Updated Time and Next Poll Time represent DataCollector time.



4. To expand or collapse a parent class, click the arrow icons next to each class name. The down arrow ▼ icon indicates that the group is expanded; the right arrow ► icon indicates that the group is collapsed.

Manage Network Views

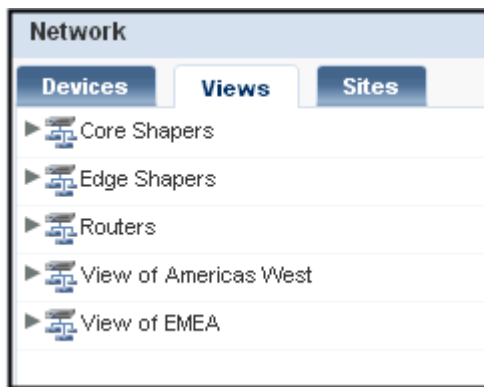
After you finish [adding devices](#) to the IntelligenceCenter network, you can define *network views*. Network views are logical groups of devices that allow you to generate reports based on specific cross-sections of traffic on your network.

You perform the following tasks when managing views:

- [create views](#)
- [modify views](#)
- [delete views](#)

Network Views Overview

Network views are logical groups of devices that allow you to generate reports based on specific cross-sections of traffic on your network. For example, you could create a network view for each business unit in your organization, allowing you to troubleshoot and monitor performance and applications for each unit independently. Creating a view is as simple as selecting devices or sub-groups from the list of network devices. You can [create network views](#) after you define your [IC network devices](#).



Create a Network View

To create a network view:

1. Select **Configure > Network**.
2. If it's not already selected, select the **Views** tab.
3. In the *Network* pane, click **Add**. The *Create Group View* dialog box appears.

Create Group View

Name *

Description

Select the Group to associate with this View:

Network *Note: You can only add devices that belong to the selected Group to this View.*

Devices

- ☐ Network
 - ☐ DataCollector1
 - ☐ Asia
 - ☐ Europe
 - ☐ North America
 - ☐ South America

4. Enter a **Name** and optionally a **Description** for the view.
5. Select a value from the **Select the Group to associate with this View** drop-down list. After you select a Group, the Sub-groups and devices contained in the group are displayed in the **Devices** box.
6. Select the devices to add to the view by selecting the box next to the device or sub-group object in the **Devices** box. If you select a sub-group, all devices in the sub-group are automatically selected.

Create Group View ?

Name * Core Shapers

Description PacketShapers - 7500 and higher

Select the Group to associate with this View:

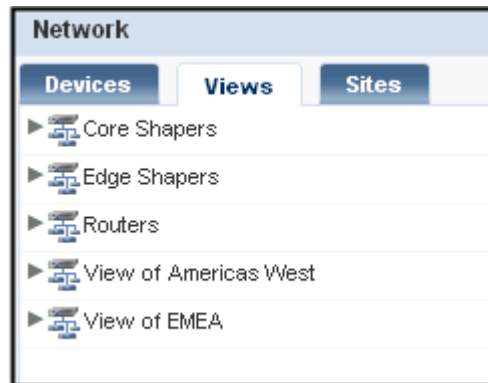
Network ▼ *Note: You can only add devices that belong to the selected Group to this View.*

Devices

- ▼ Network
 - ☐ DataCollector1
 - ☐ Asia
 - ▼ ☒ Europe
 - ☒ France
 - ▼ ☒ North America
 - ☒ East Coast
 - ☐ South America

Save Cancel

- Click **Save**. The view is added to the **Views** tab.



Modify a Network View

To modify a network view:

1. Select **Configure > Network > Views**.
2. Select the view you want to modify. The view details are displayed in the right-hand pane.

Core Shapers

General Information

Name

*

Core Shapers

Description

PacketShapers - 7500 and higher

Members

All members of this View must be a part of the same Group

▶ Europe

▼ North America

East Coast

3. To change the name of the view, enter a new value in the **Name** field.
4. To change the description of the view, enter a new value in the **Description** field.
5. To add or remove **Members**, click **Add/Edit Members** and then do the following:
 - To add a device to the view, check the checkbox that corresponds to the device or sub-group you want to add.
 - To remove devices from the view, uncheck the checkbox(es) that correspond to the sub-group or devices you want to remove.

- Click **Make Changes** to save the updated view membership.

Add / Edit View Members [X]

Select the Group to associate with this View:

Network [v] *Note: You can only add devices that belong to the selected Group to this View.*

Devices

- ▼ [icon] Network
 - ☐ [icon] DataCollector1
 - ☐ [icon] Asia
 - ☒ [icon] Europe
 - ☒ [icon] North America
 - ☐ [icon] South America


Make Changes

6. When you are done modifying the view, click **Save**.

Remove a Network View

To remove a network view that you previously created:

1. Select **Configure > Network > Views**.
2. Select the view you want to remove and click **Remove**.



3. When prompted to confirm the deletion, click **Yes**.

Manage Applications

In IntelligenceCenter (IC), an *application* is a grouping of traffic classes and/or applications that define an area of reporting interest. Application traffic will not be displayed in an IC report or portlet if there is not an application defined in IC to represent it. By default, IC includes applications that match most standard network traffic. The matching rules for each application, as related to the PacketShaper class tree, use the `/Inbound/*/<class_name>` and `/Outbound/*/<class_name>` format. With these matching rules, any standard traffic class (including most site-based traffic classes) will match the application definition as long as it is a leaf class. For example, the class `/Inbound/HTTP` and `/Inbound/Paris/HTTP` would both match the default IC application called HTTP. However, if you use custom traffic classes on your network, you will need to manually create applications or modify the default IC applications in order for the traffic for these classes to be reported in application reports and portlets.

This section describes how to manage IC [applications](#). It includes the following topics:

- [Define a new application](#)
- [Modify an existing application](#)
- [Delete an application](#)

Applications Overview

Many reports that you can run in IC detail the performance of applications on your network. Applications map to PacketShaper traffic classes, however they are not the same as traffic classes. A traffic class is logical grouping of traffic flows that share the same characteristics based on matching rules. All traffic in the traffic class must match the same rules in order to be part of the traffic class. However, you can define applications to match multiple traffic classes and/or other applications to group traffic that you are interested in reporting on as a single entity even if its flow characteristics are vastly different.


If you are simply interested in reporting on your traffic classes as they are defined on the appliances from which you're collecting data, you do not need to worry about creating applications; you can simply run traffic class reports. However, if you want to report traffic classes or traffic class groupings as applications, you must make sure that there is an application defined in IC to represent them. By default, IC includes applications that match many standard network traffic classes. The matching rules for each application, as related to the PacketShaper class tree, use the `/Inbound/*/<class_name>` and `/Outbound/*/<class_name>` format. With these matching rules, any standard traffic class (including most site-based traffic classes) will match the application definition as long as it is a leaf class. For example, the class `/Inbound/HTTP` and `/Inbound/Paris/HTTP` would both match the default IC application called HTTP.

Whenever IC discovers a new traffic class that matches one of its applications, the application becomes active. An application can only be reported on if it is (or was at one time) active. IC learns about the traffic classes on the devices it manages through its class discovery operation. By default, IC performs class discovery once a day, but this interval is [configurable](#). You can also force IC to perform a class discovery on demand.

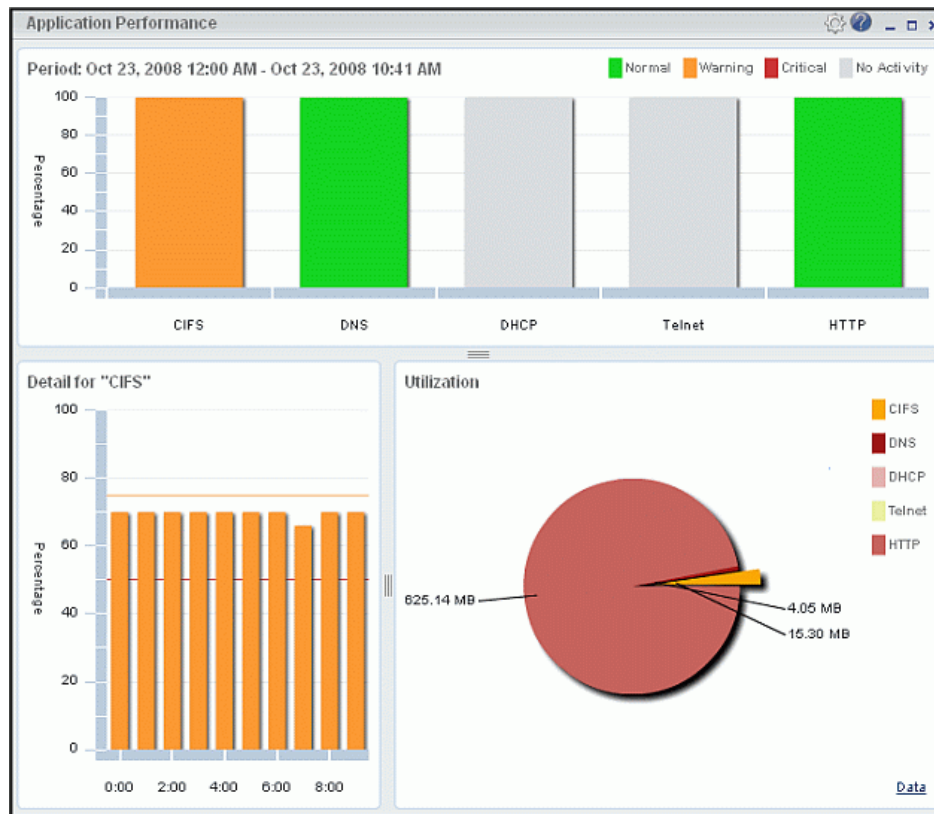
Because the default IC applications only match standard traffic classes, you will need to manually create applications for any custom traffic classes on your network that you want to report on as applications.

Applications can also contain other applications. This is useful when you want to report on a group of applications as a single entity. For example, you could create an ERP application that contains the Oracle and SAP applications. You can also [designate applications as critical applications](#). Many reports allow you to then generate reports that contain data for these critical applications only.

After you've defined your applications, you can [generate reports](#) that give you insight into which applications are consuming the most network bandwidth.



Or, you can configure [Application Performance](#) portlets that measure the performance of your applications against benchmarks you define.



Create an Application

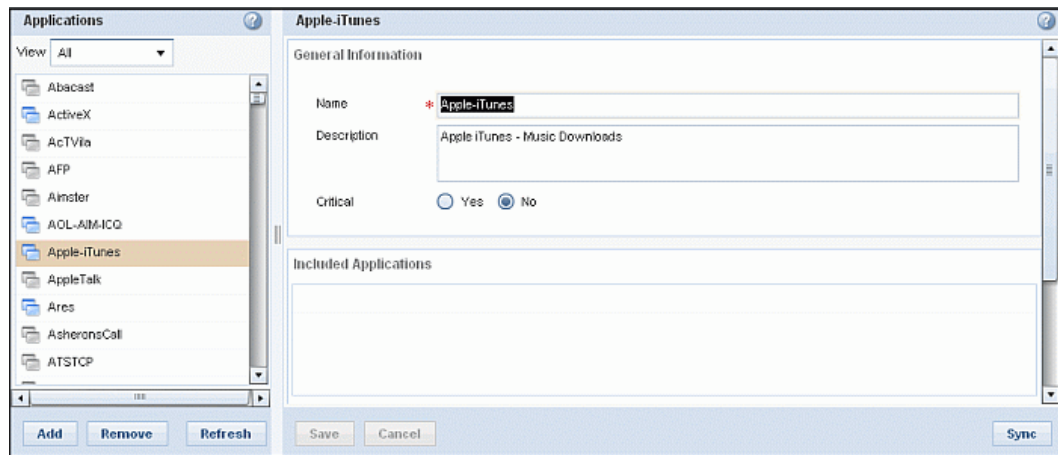
By default, IC includes applications that match most network traffic. However, there are two cases in which you might need to create an application:

- If you use custom traffic classes on your network, you will need to manually create applications or [modify existing applications](#) to match your custom classes before the traffic for these classes can be reported in application reports and portlets.
- If you want to group other applications into an area of reporting interest, you can create a hierarchical application. For example, you could create an ERP application that contains the Oracle and SAP applications.

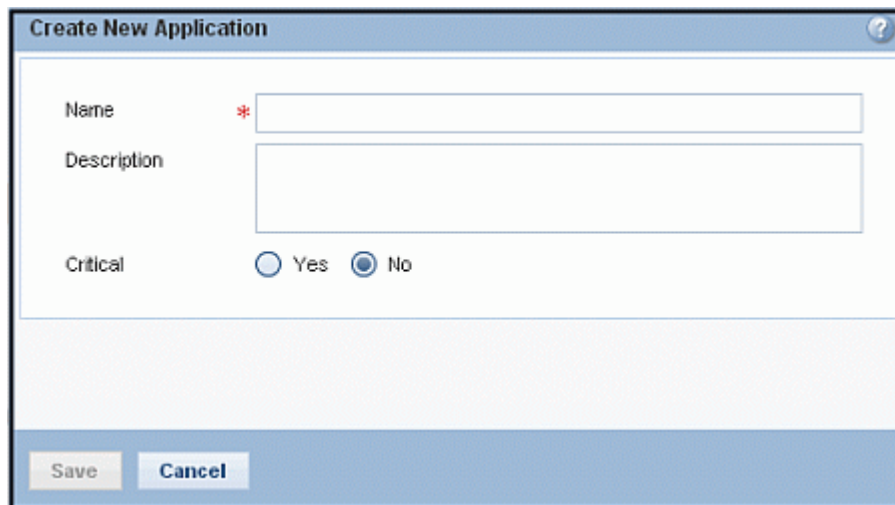
The applications you create can include any combination of applications and/or traffic class matching rules.

To create an application:


1. Select **Configure > Applications**. The application management panes are displayed. The left pane shows the list of predefined applications; if you add an application, it will be added to this list. The right pane shows the settings for the selected application.



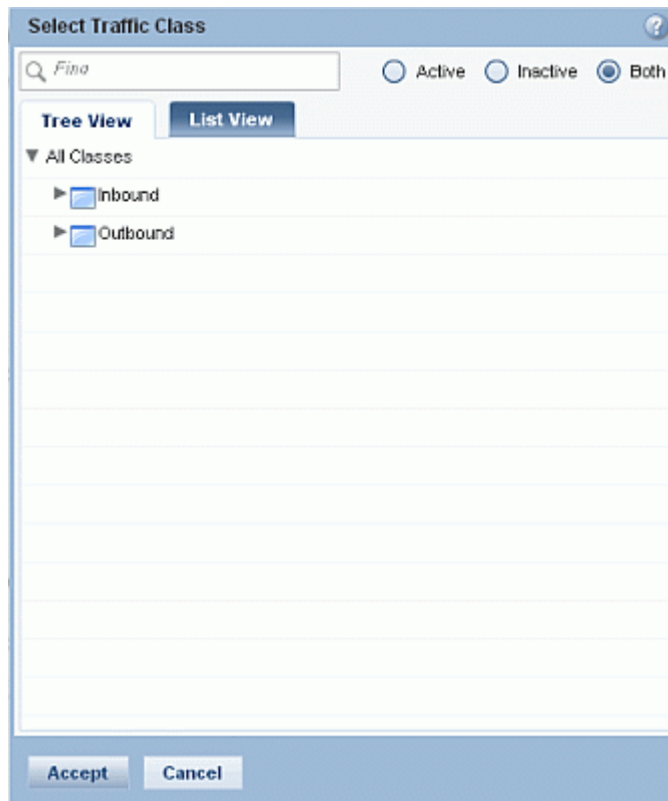
2. To create a new application, click **Add** in the *Applications* pane. The *Create New Application* dialog box is displayed.


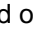


3. Enter a **Name** and optionally a **Description** of the new application.
4. If you want to designate the application as a critical application, select **Yes** in the **Critical** field. When running reports later, you will be able to specify whether to run a report against all applications or against critical applications only.
5. Click **Save**. The new application is added to the *Applications* list and a configuration window opens in the right-hand pane.
6. To define the rules that determine which traffic classes match the application, click **Add** in the **Included Traffic Classes** box expands so that you can define the traffic class matching rules for the application as follows:



- Click **Select**. The *Select Traffic Class* dialog box is displayed.



- Select the traffic class you want to match the application or type the traffic class path in the text field and click **Accept**. Note that if IC has not yet discovered the traffic class you want to add, you must [force a class discovery](#) before you can select the class (or wait for the next scheduled discovery interval). If you type the path, the path you enter may contain one or more alphanumeric strings—including the underscore (_), dash (-), or period (.)—separated by a forward slash (/). Wildcards (*) are allowed, but not at the beginning or end of the path. Additionally, the path may not end with a slash. For example, `Class.1-2`, `Class1/Class2_3`, and `Class.1/*Class2` are all valid paths.
 - Select the matching pattern from the drop-down list (**Path ends with**, **Path contains**, **Path is**, **Path starts with**).
 - Select the matching direction (**Both Directions**, **Inbound Only**, or **Outbound Only**).
 - To add another rule, repeat this step.
7. If you want to add an existing application to this application, drag it from the *Applications* pane and drop it in the **Included Applications** box.
 8. When you are done defining the application, click **Save**.
 9. (optional) Select a value from the **View** menu to filter the list of displayed applications. You can choose to display **All** applications, **Inactive** applications or **Active** applications. An active  application is an application for which a corresponding traffic class has been discovered on one of the configured data sources; an inactive  application is an application for which the corresponding traffic classes have not been discovered on the configured data sources. If DataCollector has not yet discovered traffic classes all applications will display as inactive; you must [force a class discovery](#) to see which applications are active or wait for the next discovery interval (every 24 hours by default).

Define Critical Applications

IC allows you to generate reports that detail the performance of the applications that are the most critical to your business. Before you can generate reports on your critical applications, you must first designate which applications to include as critical.

To designate an application as critical:

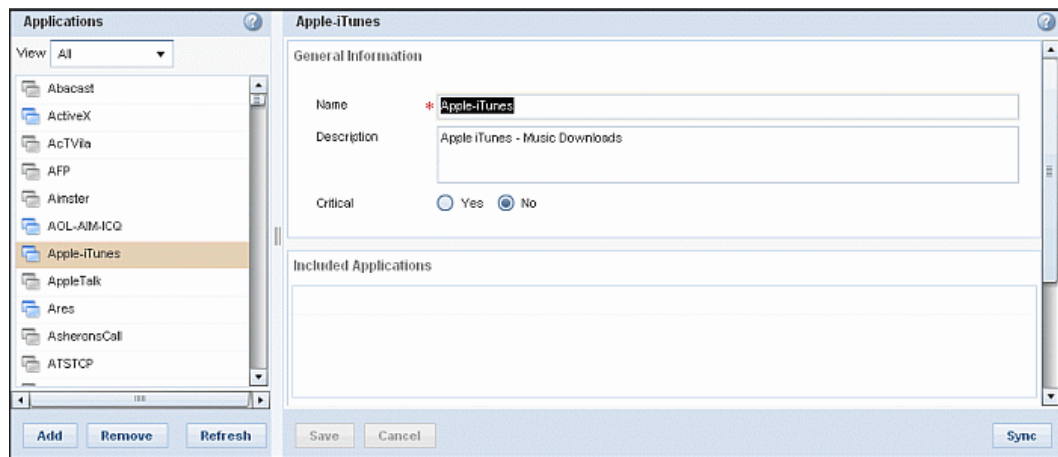
1. Select **Configure > Applications**.
2. Select the application you want to designate as a critical application in the Applications pane.
3. Set the **Critical** radio button to **Yes**.
4. Click **Save**.

Modify an Application

By default, IC includes applications that match most standard network traffic. The matching rules for each application, as related to the PacketShaper class tree, use the `/Inbound/*/<class_name>` and `/Outbound/*/<class_name>` format. With these matching rules, any standard traffic class (including most site-based traffic classes) will match the application definition as long as it is a leaf class. However, if you want a traffic class that does not match this pattern to match an existing application, you must modify the application to include rules that will match the traffic class. You can also add an existing application to the application, creating a hierarchical applications.

To modify an application definition:

1. Select **Configure > Applications**.
2. In the *Applications* pane, select the application you want to modify. The application details are displayed in the right pane.

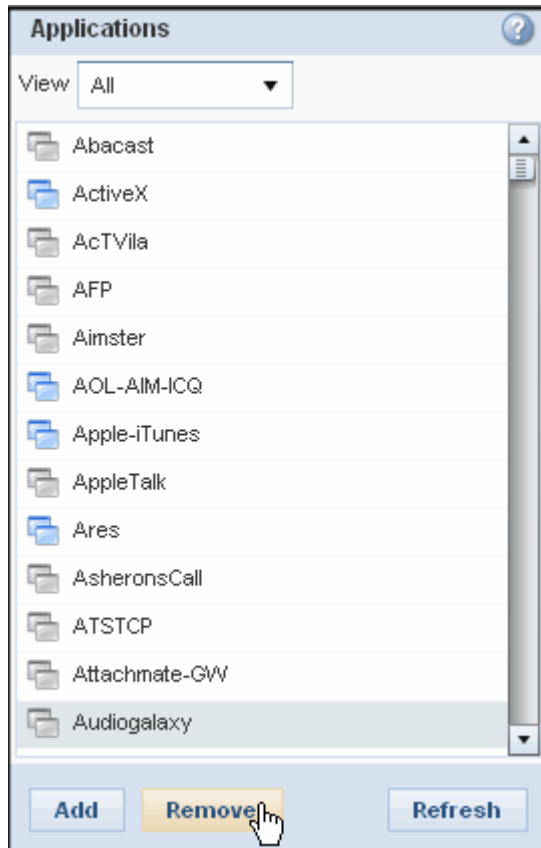


3. Modify the **Name** and **Description** fields as desired.
4. Add, modify, or remove traffic class matching rules as follows.
 - To add a traffic class matching rule, click **Add**. The **Included Traffic Classes** box expands. To select the traffic classes to match, click **Select**. Select the traffic class from the Select Traffic Class dialog box and then click **Accept**. Then select the matching pattern (**Path ends with**, **Path contains**, **Path is**, **Path starts with**) and the matching direction (**Both Directions**, **Inbound Only**, or **Outbound Only**). Note that if IC has not yet discovered the traffic class, you will not yet be able to select it. In this case, you must [force a class discovery](#) before adding the new rule.
 - To modify an existing matching rule, click the arrow icon to expand the rule. You can then modify the traffic class to select, the matching pattern, or the direction.
 - To remove a matching rule, click the corresponding remove icon.
5. Add or remove an existing application from this application as follows:
 - To add an existing application, drag it from the *Applications* pane and drop it in the **Included Applications** box.
 - To remove an application, click the corresponding **Remove** link in the **Included Applications** box. When prompted, confirm that you want to remove the entry.
6. When you are done modifying the application, click **Save**.

Remove an Application

If there is an application that you do not use or that you do not want to report on, you can remove it from IntelligenceCenter. To remove an application:

1. Select **Configure > Applications**.
2. In the *Applications* pane, select the application you want to delete and click **Remove**.



3. When prompted to confirm the deletion, click **Yes**.

Configure Static Reporting

If you have configured data collection, you can generate reports that aggregate the metric (ME) and flow detail record (FDR) data that is collected by the devices throughout your network. This allows you to compare, correlate, and summarize network behavior throughout your organization.

This section describes the [collection of static reports](#) and describes how to manage reporting, including:


- [Running a report](#)
- [Searching for and viewing a report](#)
- [Deleting a report](#)

The static reports that you run can either be generated immediately or you can schedule them to run some time in the future, either once or on a recurring basis. IntelligenceCenter provides five predefined report schedules that you can choose from. You can [modify the predefined schedules](#) or you can [create your own schedule definitions](#). Note that the specific [reports](#) you can run depends on whether you are collecting ME data, FDR data, or both. Additionally, if you are collecting both types of data and you want the ME data and FDR data to correlate, you must collect both types of data from the same set of data sources (although because FDR is delivered over UDP, which is not guaranteed, there may still be slight differences between ME and FDR).

Keep in mind that the time range you select when running a report determines what [database table](#) DataCollector will pull the data from. Therefore, when [selecting a time range](#), you must consider whether DataCollector has had a chance to roll data up into the corresponding database table. If the corresponding database table does not contain any data yet, DC will look for the data in tables with more granularity, which may cause delays in generating the report. In addition, the more data you have, the longer your reports will take to run. If the reports you run regularly are taking a long time to generate, consider [scheduling them](#) instead of running them on demand.

Reporting Overview

IntelligenceCenter provides a comprehensive [set of reports](#) that allow you to monitor the performance of the applications, devices, and hosts that are running on your network. The data that is displayed in the IntelligenceCenter reports is collected from the PacketShaper appliances and other network devices, such as routers, switches, or bridges, throughout your network and aggregated according to options you define when you [run the report](#). You can generate reports that detail performance across an entire network group or sub-group, or you can define a specific device, [site, geographic region or logical group](#)—such as a business unit—on which to report. These reports enable you to quickly and easily identify emerging network trends, accelerate network response times and troubleshooting, and perform capacity planning.



With IntelligenceCenter reporting you can:

- View high-level information about your network and then drill-down into the information that you are interested in.
- Report in any time zone. All data that is collected is stored in Coordinated Universal Time (UTC). When you run a report, you select the specific time zone for which you want to view data and IntelligenceCenter does the work of normalizing the data and presenting it.
- Create [consolidated reports](#) that allow you to configure report options for a sequence of reports—both top-level reports as well as drill-down reports, from one or more network groups—and schedule them to run together and be output as a single report.
- Use predefined time ranges or define the specific start and end dates and times that are of interest to you.
- Define a set of [business hours](#) that correspond to the hours when your network usage is most active (that is, the hours when people are working). You can then run your reports to show data for only those hours. This way, statistical information displayed in the reports is not skewed by the hours when there is very little traffic, such as weekends and evenings.
- Run reports ad-hoc or on a schedule.
- Save reports in the report archive and/or in the file system for viewing by multiple IntelligenceCenter users.
- View and print reports, either PDF, HTML, Microsoft Word (DOC), or Microsoft PowerPoint (PPT).
- Filter the report archive by time range or keyword to locate the report you want to view.
- Email PDFs of generated reports to a specified set of recipients.

To generate reports, you must have a valid DataCollector [license](#) and you must have installed and [configured DataCollector](#) to collect data from your PacketShaper appliances and/or other network devices. Some reports are based on FDR data, some on ME data, and some use both ME and FDR data, therefore, the reports that you can view depend on which type of data you are collecting. For best correlation between reports, you should collect both ME and FDR data from the same set of devices. Keep in mind, however, that because FDR data is sent over UDP,

delivery is not guaranteed and there may therefore be some slight differences between the ME and FDR reports.

The time range you select when running a report determines what [database table](#) DataCollector will pull the data from. Therefore, when [selecting a time range](#), you must consider whether DataCollector has had a chance to roll data up into the corresponding database table. If the corresponding database table does not contain any data yet, DC will look for the data in tables with more granularity, which may cause delays in generating the report. In addition, the more data you have, the longer your reports will take to run. If the reports you run regularly are taking a long time to generate, consider [scheduling them](#) instead of running them on demand.

Time Synchronization Overview

IntelligenceCenter allows you to generate reports that aggregate data from network devices throughout your network so that you can compare, correlate, and summarize network behavior throughout your organization. For the report data to be meaningful, all of the devices from which you collect data must have their clocks synchronized. Therefore, each appliance must have Simple Network Time Protocol (SNTP) enabled. When SNTP is enabled on a PacketShaper appliance, the appliance clocks are synchronized with a Network Time Protocol (NTP) server.

In addition, when DataCollector collects data from the appliances throughout your network, it converts the time stamp on the data from the local time of the appliance to Coordinated Universal Time (UTC). This way of normalizing the data allows you to generate aggregated reports that present statistics representing the conditions of your network as a whole at a specific point in time.

Therefore, in order for DataCollector and your PacketShaper appliances to be able to reconcile data accurately, DataCollector and the appliances from which it collects must all be configured to get their date and time automatically from an NTP server. For DataCollector, you must specify the NTP server URL during installation. For instructions on how to enable SNTP on your PacketShaper appliances, refer to [Modify Date and Time Settings](#) in PacketGuide. For best results, consider using DC as the time server for your PacketShaper appliances.

Reporting Time Ranges and Data Granularity

The time range you select when running a report or portlet determines what database table DataCollector will pull the data from. Therefore, when selecting a time range, you must consider whether DataCollector has had a chance to roll data up into the corresponding database table.

The following table lists the possible time ranges and the corresponding data granularity that is used for each range.

Selected Time Range	Database Table	Data Granularity
< 15 minutes	Raw	raw
< 1 hour	Hourly	15 minutes
< 1 day	Hourly	15 minutes
< 1 month	Daily	1 hour
< 1 year	Monthly	1 day
> 1 year	Yearly	1 month

If the corresponding database table does not contain any data yet, one of the following will happen:

- For portlets, the portlet will not display any data. For example, if you select a seven day time range, IntelligenceCenter will pull the data from the DataCollector's daily table. However, if you do not yet have any data rolled up into the daily table, the portlet cannot display any data.
- For reports, DC will look for the data in tables with more granularity, which may cause delays in generating the report. For example, if you select a seven-day time span, IC will look for the data in the daily table. However, if you do not yet have any data rolled up into the daily table, DC will have to get the data from the hourly table.

You should also keep in mind that because DC rolls data up over time, older data may not contain as much granularity as newer data (depending on what [data retention policies](#) you have configured). Therefore, if you select a short time span on older data, DC may not have retained the data and therefore will not be able to display it in the report.

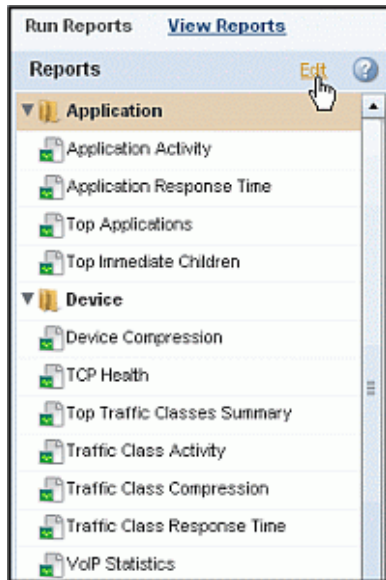
When running a report you should also consider the current date or time in relation to the report time range you select. For example, if you run a report using the This month time range on the first day of the month, the report will not return any data, because the This month time range reports from midnight on the first day of the month up through the last full day; if a full day has not yet completed there will not be any data to display.

Customize the Report List

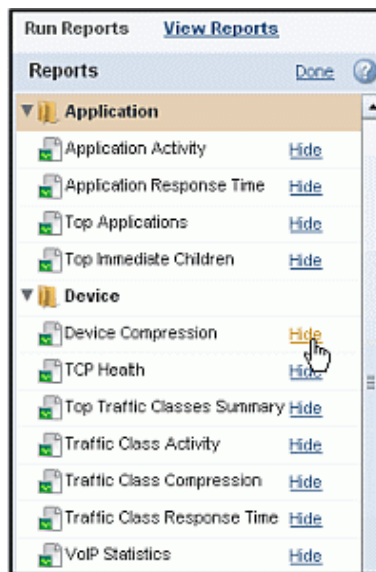
IntelligenceCenter includes a large variety of reports that you can run. However, if you find that you do not use some of the reports provided in the application, you may want to customize which reports are displayed.

To customize the report list:

1. Select **Report > Run Reports**.
2. Click the **Edit** link in the **Reports** pane. A **Hide** link appears next to each report name.



3. To remove a report from the **Reports** pane, click the corresponding **Hide** link (or to add back a report that you previously removed, click **Show**).



4. When you are done modifying the list of reports, click the **Done** link at the top of the **Reports** pane.

Run a Report

There are three ways to run a static report:

- [Run a report for immediate viewing only](#)
- [Run a report for immediate viewing and save it in the report archive](#)
- [Schedule a report to run sometime in the future](#)

Run and View a Static Report

To run a report for immediate viewing only:

1. Click the **Report** tab.
2. Select the [report](#) you want to run from the *Reports* pane. When you select a report, a *Report Configuration* pane opens.

Report Configuration

Top Applications
Shows the top bandwidth-consuming applications

Time Range: Today ☐ Consider Business Hours

Formatting

View Format: ☒ HTML ☐ PDF ☐ DOC ☐ PPT

Run Options: ☐ Run and Archive ☒ Run and View ☐ Run on a Schedule

Report Options

Network Group or Device *

Critical Application ☐ Yes ☒ No

Max Count *

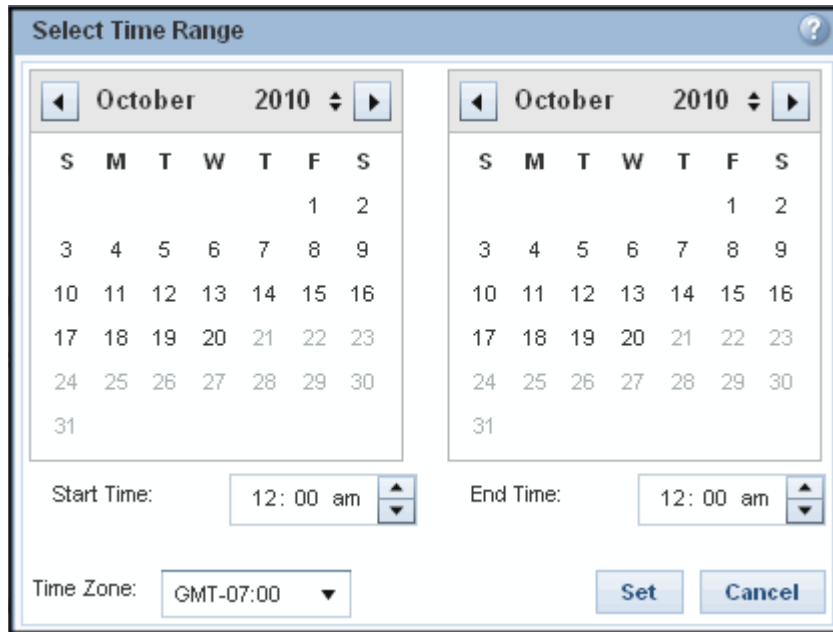
Sort Column

Sort Direction ☒ Descending ☐ Ascending

Email To [Add](#)

3. Select **Run and View**.
4. To set the time range for the report, select a [predefined time range](#) from the **Time Range** drop-down list or select **Other** and then define a custom time span using the *Select Time Range* dialog box.

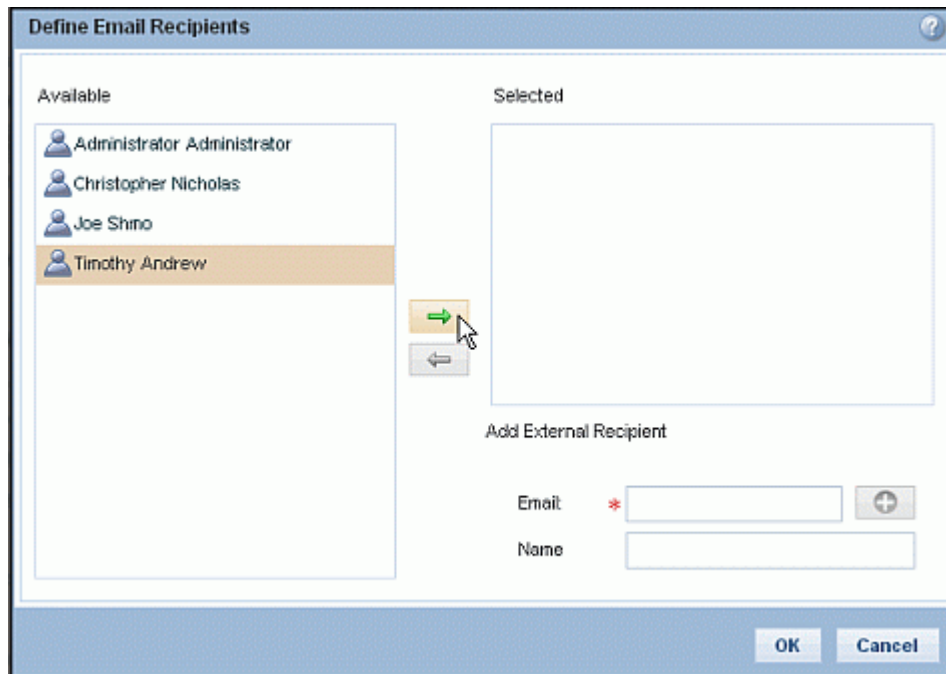
Note: For best performance, you should use a predefined time range whenever possible.





The dialog box titled "Select Time Range" contains two side-by-side calendar controls for October 2010. Each calendar has a header with navigation arrows and the month/year. The days of the week are abbreviated as S, M, T, W, T, F, S. The dates are displayed in a grid. Below the calendars are two time selection fields: "Start Time:" and "End Time:", each with a text box showing "12: 00 am" and up/down arrows. At the bottom left is a "Time Zone:" dropdown menu showing "GMT-07:00". At the bottom right are "Set" and "Cancel" buttons.

To define the custom time span:

- Set the start and end date by selecting values from the calendars. The left-hand calendar allows you to set the start date; the right-hand calendar allows you to set the end date. The end date cannot exceed the current date.
 - Set the **Start Time** and **End Time** using the up and down arrows in each field or by entering the desired time in the hours and minutes fields.
 - If you want to display report data for a time zone other than the time zone where your client system is located (the default time zone), select a new value from the **Time Zone** field. All data that is collected is stored in Coordinated Universal Time (UTC) and can be normalized and rendered for any time zone. Time zone values are displayed as Greenwich Mean Time (GMT) offsets; that is, how many hours ahead (+) or behind (-) GMT the time is.
 - When you are done setting the time span, click **Set**.
5. If you want the report to show only data for the hours designated as business hours, select the **Consider Business Hours** checkbox. Note that this option only applies if you have [set up business hours reporting](#) for the DataCollector that will report data for the network group or device you select when you configure the report options.
 6. Specify whether you want the report to be generated as **HTML**, **PDF** (Adobe Portable Document Format), **DOC** (Microsoft Word document format), or **PPT** (Microsoft PowerPoint format).
 7. Specify the other *Report Options* you want to use to generate the report. The specific report options depend on which report you are running. The report options for each report are detailed in the Reference section at the end of this guide.
 8. If you want to email a PDF of the generated report to one or more recipients (and you have [enabled support for email notification](#)), click **Add** in the *Email To* section. The *Define Email Recipients* dialog box is displayed.



9. Add email recipients as follows:
 - If you want to email the report to an existing IntelligenceCenter user, select the user from the **Available** column and click the  button to move the user to the **Selected** column. To select multiple, consecutive users Shift+Click. To add multiple, non-consecutive users, Ctrl+Click.
 - If you want to email the report to a recipient who does not have an IntelligenceCenter [user profile](#), enter the user's **Email** address and optionally the user's **Name** and then click the add  button. The recipient is added to the **Selected** list. Repeat this step for each user you want to add. You can add any email address that is recognized by your SMTP server, including group alias addresses. Note, however, that any external users you add are not saved to the IntelligenceCenter database and you will have to readd them each time you run a report that you want to send to that person. Therefore, if you find that you are sending reports to the same people frequently, you should [create a user profile](#) for the person.
 - When you are done defining the recipients, click **OK**. The recipients you added are displayed on in the *Email To* section of the *Report Configuration* pane. Note that you can delete a recipient by clicking the **Remove** link next to the recipient's user profile name or email address.
10. Click **Run**. IntelligenceCenter runs the report and then displays it in your default browser. Additionally, a PDF of the report is emailed to any recipients you defined.

Run and Archive a Static Report

To run a report both for immediate viewing and for saving in the report archive:

1. Click the **Report** tab.
2. Select the [report](#) you want to run from the *Reports* pane. When you select a report, a *Report Configuration* pane opens.

Report Configuration

Top Applications
Shows the top bandwidth-consuming applications

Time Range: Today ☐ Consider Business Hours

Formatting

View Format: ☒ HTML ☐ PDF ☐ DOC ☐ PPT

Run Options: ☒ Run and Archive ☐ Run and View ☐ Run on a Schedule

Report Name:

Report Options

Network Group or Device:

Critical Application: ☐ Yes ☒ No

Max Count:

Sort Column:

Sort Direction: ☒ Descending ☐ Ascending

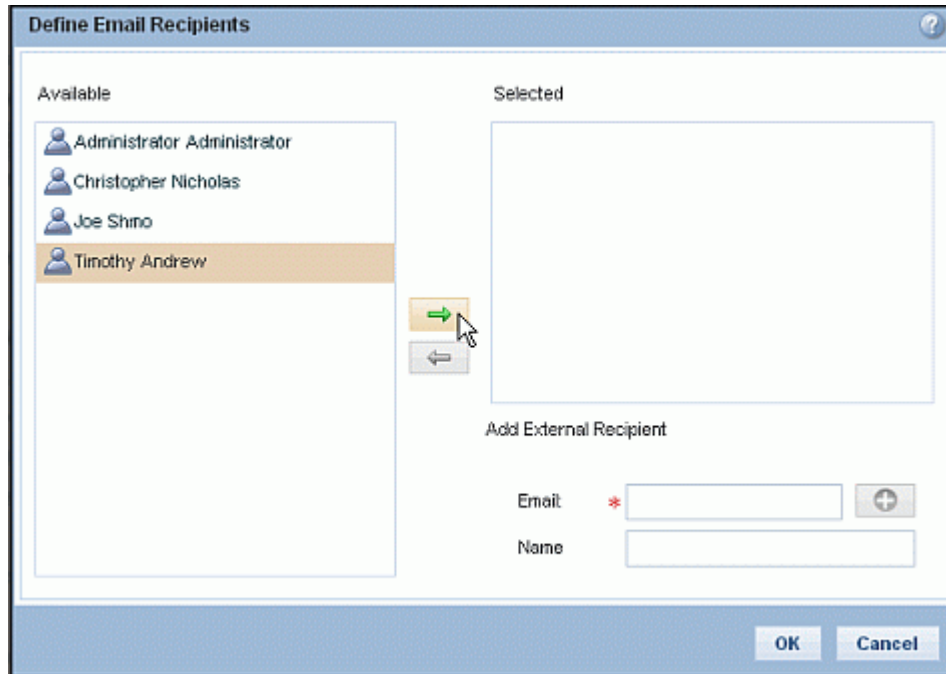
Email To: [Add](#)

3. Select **Run and Archive**.
4. Specify a **Report Name**. Because this is the name that will identify the report after it runs, you should try to be descriptive so that you'll be able to [find the report](#) in the *Archived Reports* pane.
5. To set the time range for the report, select a [predefined time range](#) from the **Time Range** drop-down list or select **Other** and then define a custom time span using the *Select Time Range* dialog box.



Note: For best performance, you should use a predefined time range whenever possible.

To define the custom time range:

- Set the start and end date by selecting values from the calendars. The left-hand calendar allows you to set the start date; the right-hand calendar allows you to set the end date. Note that the end date cannot exceed the current date.
 - Set the **Start Time** and **End Time** using the up and down arrows in each field or by entering the desired time in the hours and minutes fields.
 - If you want to display report data for a time zone other than the time zone where your client system is located (the default time zone), select a new value from the **Time Zone** field. All data that is collected is stored in Coordinated Universal Time (UTC) and can be normalized and rendered for any time zone. Time zone values are displayed as Greenwich Mean Time (GMT) offsets; that is, how many hours ahead (+) or behind (-) GMT the time is.
 - When you are done setting the time span, click **Set**.
6. If you want the report to show only data for the hours designated as business hours, select the **Consider Business Hours** checkbox. Note that this option only applies if you have [set up business hours reporting](#) for the DataCollector that will report data for the network group or device you select when you configure the report options.
 7. Specify whether you want the report to be generated as **HTML**, **PDF** (Adobe Portable Document Format), **DOC** (Microsoft Word document format), or **PPT** (Microsoft PowerPoint format).
 8. Specify the other *Report Options* you want to use to generate the report. The specific report options depend on which report you are running. The report options for each report are detailed in the Reference section at the end of this guide.
 9. If you want to email a PDF of the generated report to one or more recipients (and you have [enabled support for email notification](#)), click **Add** in the *Email To* section. The *Define Email Recipients* dialog box is displayed.



10. Add email recipients as follows:

- If you want to email the report to an existing IntelligenceCenter user, select the user from the **Available** column and click the  button to move the user to the **Selected** column. To select multiple, consecutive users Shift+Click. To add multiple, non-consecutive users, Ctrl+Click.
- If you want to email the report to a recipient who does not have an IntelligenceCenter [user profile](#), enter the user's **Email** address and optionally the user's **Name** and then click the add  button. The recipient is added to the **Selected** list. Repeat this step for each user you want to add. You can add any email address that is recognized by your SMTP server, including group alias addresses. Note, however, that any external users you add are not saved to the IntelligenceCenter database and you will have to readd them each time you run a report that you want to send to that person. Therefore, if you find that you are sending reports to the same people frequently, you should [create a user profile](#) for the person.
- When you are done defining the recipients, click **OK**. The recipients you added are displayed on in the *Email To* section of the *Report Configuration* pane. Note that you can delete a recipient by clicking the **Remove** link next to the recipient's user profile name or email address.

11. Click **Run**. IntelligenceCenter runs the report and then displays it in your default browser. Additionally, a PDF of the report is emailed to any recipients you defined. You can [view reports](#) in the *Archived Reports* pane at any time.

Schedule a Static Report

The report scheduling feature allows you to schedule a report to run at some time in the future—either once or on a continuing basis.

Note: You must [schedule custom reports](#) from the **Report > Custom Reports** tab.

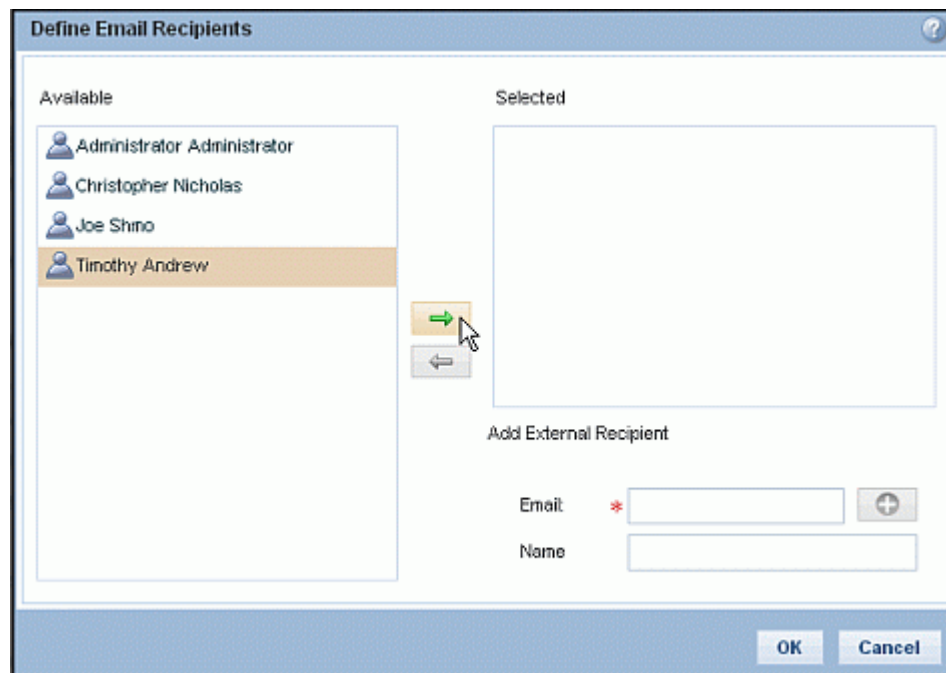
To schedule a report:

1. Click the **Report** tab.
2. Select the [report](#) you want to run from the *Reports* pane. When you select a report, a *Report Configuration* pane opens.

3. Select the **View Format** for the report: **HTML**, **PDF** (Adobe Portable Document Format), **DOC** (Microsoft Word document format), or **PPT** (Microsoft PowerPoint format).
4. Select **Run on a Schedule**. When you select this option, additional scheduling options are displayed.



Select Schedule	System Daily ▼	
	Start Time:	Jun 7, 2011 12:00 PM
	End Time:	Never end
	Recurring:	Every day
Time Range	Year ▼	<input type="checkbox"/> Consider Business Hours

5. In the **Select Schedule** field, select the schedule you want to use to run this report, or select **Create Schedule** to [define a custom schedule](#) for this report.
6. Pick a time range for the reports to be generated by selecting a value from the **Time Range** drop-down list (**Year**, **Month**, **Week**, or **Day** or a **Custom** number of days). Scheduled reports use the start time that the report is scheduled to run as the end time of the report. The start time is then determined by subtracting the time span from the end time. For example, if a report is scheduled to run today at 10 AM for a **Time Range** of **Month**, the start time will be 10 AM today minus one month.
7. If you want the report to show only data for the hours designated as business hours, select the **Consider Business Hours** checkbox. Note that this option only applies if you have [set up business hours reporting](#) for the DataCollector that will report data for the network group or device you select when you configure the report options.
8. Specify the other *Report Options* you want to use to generate the report. The specific report options depend on which report you are running. The report options for each report are detailed in the Reference section at the end of this guide.
9. If you want to email a PDF of the generated report to one or more recipients (and you have [enabled support for email notification](#)), click **Add** in the *Email To* section. The *Define Email Recipients* dialog box is displayed.



The **Define Email Recipients** dialog box is shown. It features two main panes: **Available** and **Selected**. The **Available** pane contains a list of users: Administrator Administrator, Christopher Nicholas, Joe Shino, and Timothy Andrew. Timothy Andrew is currently selected. Between the panes are two arrow buttons: a green right-pointing arrow and a grey left-pointing arrow. The **Selected** pane is currently empty. Below the panes is an **Add External Recipient** section with fields for **Email** (marked with a red asterisk) and **Name**, each followed by a plus icon. At the bottom right are **OK** and **Cancel** buttons.

10. Add email recipients as follows:

- If you want to email the report to an existing IntelligenceCenter user, select the user from the **Available** column and click the  button to move the user to the **Selected** column. To select multiple, consecutive users Shift+Click. To add multiple, non-consecutive users, Ctrl+Click.
- If you want to email the report to a recipient who does not have an IntelligenceCenter [user profile](#), enter the user's **Email** address and optionally the user's **Name** and then click the add  button. The recipient is added to the **Selected** list. Repeat this step for each user you want to add. You can add any email address that is recognized by your SMTP server, including group alias addresses. Note, however, that any external users you add are not saved to the IntelligenceCenter database and you will have to readd them each time you run a report that you want to send to that person. Therefore, if you find that you are sending reports to the same people frequently, you should [create a user profile](#) for the person.
- When you are done defining the recipients, click **OK**. The recipients you added are displayed on in the *Email To* section of the *Report Configuration* pane. Note that you can delete a recipient by clicking the **Remove** link next to the recipient's user profile name or email address.

11. Click **Run**. An entry for the report will be added to the [Scheduled Tasks](#) list. After the report runs, an entry will appear on the *Archived Reports* pane. You can [view reports](#) in the *Archived Reports* pane at any time.

Predefined Report Time Ranges

Each static report provides a set of predefined time ranges that have been designed to optimize the speed at which the report runs. To optimize performance, the predefined report time ranges are rounded to the data point that makes the most sense for the selected range based on the various database tables from which DC pulls data. For example, the **Last quarter** report is rounded to the last full quarter; the **Today** report is rounded to the last full hour; and the **This month** report is rounded to the last full day. In addition, the actual time range that is used to report data depends on the date and time at which the report runs (and the actual time range is displayed on the report header).

Note: Sometimes if you run an FDR report using the **Last hour** time range very close to the start of a new hour, DC may not have finished rolling data up into the previous hour's table in the database and the report will return with no data. If this happens, try running the report again in a few minutes.

You can select one of the following time ranges when running a static report:




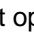
Time Range	Description
Last hour	Reports data for the last full hour. For example, if the report runs at 11:20, this selection would show data collected between 10:00:00 and 10:59:59. This is the default time range.
Today	Reports data from midnight today up through last full hour. For example, if the report runs at 11:20, the report will display data collected between 00:00:00 and 10:59:00 today.
Yesterday	Reports data from 00:00:00 to 23:59:59 yesterday.
Last seven days	Reports data for the previous seven full days. For example, if the report runs at 11:20 on 11/11, the report would show data from 00:00:00 on 11/4 through 23:59:59 on 11/10.
This week	Reports data from 00:00:00 on Monday through the last full hour today. For example, if the report runs at 11:20 on 11/11 (a Tuesday), the report would show data from 00:00:00 on Monday 10/10 through 10:59:59 on 11/11.
Last week	Reports data from 00:00:00 on Sunday of the previous week through 23:59:59 last Saturday. For example, if you ran the report on 11/11, it would show data from 00:00:00 Sunday, 11/2 through 23:59:59 on Saturday, 11/8.
This month	Reports data from 00:00:00 on the first day of the current month through the last full day. For example, if the report ran on 11/11, it would show data from 00:00:00 on 11/1 through 23:59:59 on 11/10.
Last month	Reports data from 00:00:00 on the first day of the last month through 23:59:59 on the last day of last month.
Last quarter	Reports data for the last full quarter based on the calendar year (where quarter 1 is January through March; quarter 2 is April through June; quarter 3 is July through September; and quarter 4 is October through December). For example, if the report ran on 11/11, it would show data from 00:00:00 on July 1st through 23:59:59 on September 30th.

Find a Report

By default the *Archived Reports* pane displays all reports that were generated this week. You can change how the reports are filtered or grouped or search for a specific report as follows:

1. Select **Report > View Reports**. The reports that have been run and archived are displayed.
2. To filter the reports that are displayed:
 - Select a value from the **From** field drop-down list. You can choose to display reports that were generated **Today**, **This Week** (the default), **Last Week**, **This Month**, **Last Month**, **This Quarter**, or **Last Quarter** or you can select **Other** and then click **Start Time** and/or **End Time** and select a date from the pop-up calendar(s).
 - If you want to display reports with a specific string in the name, enter the string in the text box.
 - Click **Find**. The list of archived reports is updated to display only those reports that match the criteria you entered.
3. To change how the displayed reports are grouped:
 - Select a value from the **Arrange By** drop-down list. You can choose to group the reports by **Created On** (the default), **Report** (for example, Top Applications), or **Category** (for example, Application, Device, Host).
 - If there are a large number of reports, you may want to collapse or expand a group of reports by clicking the arrow icons next to the group name (for example, if you arranged the reports by **Category**, you might want to collapse the Device and Application groups so that only the Host reports are visible). The down arrow ▼ icon indicates that the group is expanded; the right arrow ► icon indicates that the group is collapsed.
4. To sort the order in which the reports are displayed within each grouping, click the column header by which to sort. For example, click the **Start Time** column header to sort the archived reports in each group according to the starting date and time of the report time span. When you sort the report archive by a particular column, an arrow appears next to the column header. An up arrow ▲ icon indicates that the reports are being sorted in ascending order and a down arrow ▼ icon indicates that the reports are being sorted in descending order. To switch the order in which the reports are sorted (ascending or descending), click the column header again.

Name	▲ Start Time	End Time	Created On	Open As
------	--------------	----------	------------	---------

5. To view details about a report without opening it, hover over the report entry.
6. After you have found the report you are looking for, you can view it by clicking the icon that corresponds to the application in which you want to view the report in the **Open As** column of the report row. If you select , the report opens in your default browser. If you click , the report opens in Adobe Reader. If you select , the report opens in Microsoft Word. If you click , the report opens in Microsoft PowerPoint. Note that custom reports are only available in PDF or HTML.
7. To clear your search filters and display all archived reports for the selected time frame, click **Clear**.

From	This Week ▼	class ▼	Find	Clear 
------	-------------	---------	------	---




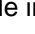
View a Report

When you [run a report](#), you can decide whether to [view it immediately](#), [save it to the report archive](#), or [schedule it to run sometime in the future](#). Whenever you run and save a report or a scheduled report runs (including the [consolidated reports](#) you schedule), an entry for the report is added to the *Archived Reports* pane on the **Report** section of IntelligenceCenter.

To view a report:

1. Select **Report > View Reports**. The reports that have been run and archived are displayed.
2. To filter the reports that are displayed:
 - Select a value from the **From** field drop-down list. You can choose to display reports that were generated **Today**, **This Week** (the default), **Last Week**, **This Month**, **Last Month**, **This Quarter**, or **Last Quarter** or you can select **Other** and then click **Start Time** and/or **End Time** and select a date from the pop-up calendar(s).
 - If you want to display reports with a specific string in the name, enter the string in the text box.
 - Click **Find**. The list of *Archived Reports* pane is updated to display only those reports that match the criteria you entered.
3. To change how the displayed reports are grouped:
 - Select a value from the **Arrange By** drop-down list at the bottom of the *Archived Reports* pane. You can choose to group the reports by **Created On** (the default), **Report** (for example, Top Applications, VoIP Statistics), or **Category** (for example, Application, Device, Host).
 - If there are a large number of reports, you may want to collapse or expand a group of reports by clicking the arrow icons next to the group name (for example, if you arranged the reports by **Category**, you might want to collapse the Device and Application groups so that only the Host reports are visible). The down arrow ▼ icon indicates that the group is expanded; the right arrow ► icon indicates that the group is collapsed.
4. To sort the order in which the reports are displayed within each grouping, click the column header by which to sort. For example, click the **Start Time** column header to sort the archived reports in each group according to the starting date and time of the report time span. When you sort the report archive by a particular column, an arrow appears next to the column header. An up arrow ▲ icon indicates that the reports are being sorted in ascending order and a down arrow ▼ icon indicates that the reports are being sorted in descending order. To switch the order in which the reports are sorted (ascending or descending), click the column header again.

Name	▲ Start Time	End Time	Created On	Open As
------	--------------	----------	------------	---------

5. To view details about a report without opening it, hover over the report entry.
6. After you have found the report you are looking for, you can view it by clicking the icon that corresponds to the application in which you want to view the report in the **Open As** column of the report row. If you select , the report opens in your default browser. If you click , the report opens in Adobe Reader. If you select , the report opens in Microsoft Word. If you click , the report opens in Microsoft PowerPoint.

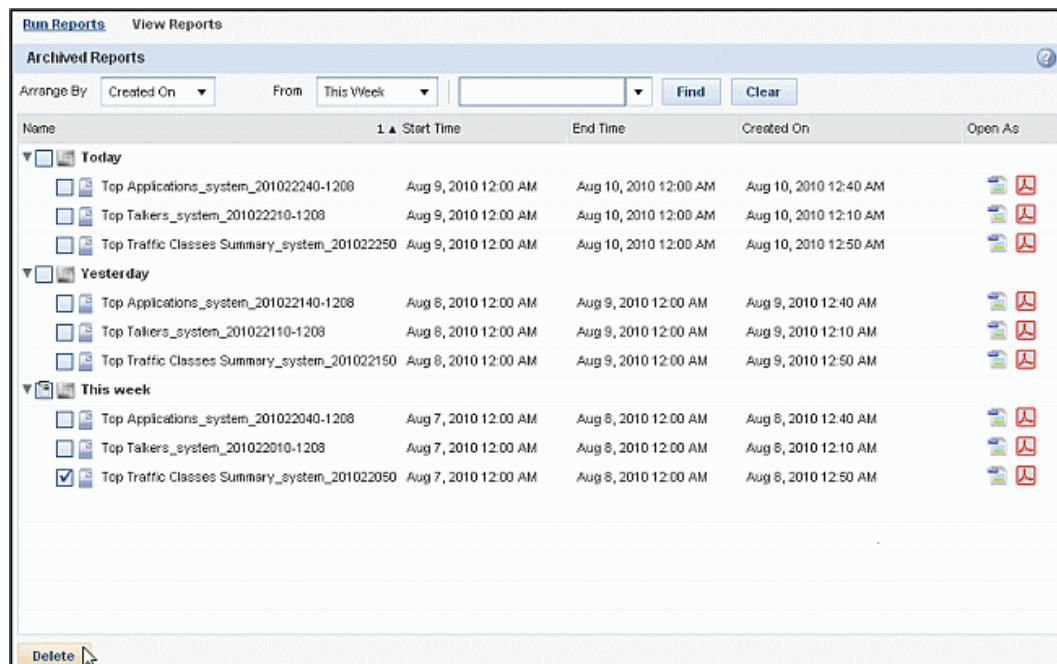
Note: Custom reports are only available in PDF or HTML.

Delete a Report

Whenever you run and save a report or a scheduled report runs, an entry for the report is added to the *Archived Reports* pane in the **Report > View Reports** section of IntelligenceCenter.

To delete a report:

1. Click the **Report** tab.
2. [Locate the report you want to delete.](#)
3. To delete the report or group of reports, select the checkbox next to the report entry and then click the **Delete** button at the bottom of the *Archived Reports* pane. To delete an entire group of reports, click the checkbox next to the group heading (for example, if you arranged the reports by **Category**, you could select the checkbox next to the Host grouping to delete all Host reports).



Note: If you want to view details about a report before deleting it, just hover over the report entry. Or, you can view the report by clicking (PDF), (HTML), (DOC), or (PPT) in the **Open As** column of the report row.

Define Report Schedules

With the scheduling feature, you can schedule reports to be run at a specific time, on a specific day, on an ongoing basis or only once. IntelligenceCenter provides five system defined schedules: **System Hourly**, **System Daily**, **System Weekly**, **System Monthly**, and **System Yearly**. You can modify the specific recurrence patterns for these predefined schedules. For example, for the daily schedule, you can specify at what time of day you want the report to run. Similarly, for a monthly schedule you can define the specific day of the month on which to run the report. Or, if you want to run a report on a schedule other than one of the system defined schedules — such as biweekly, bimonthly, or quarterly — you can define your own custom report schedules. You can:

- [Modify a schedule definition](#)
- [Add a new schedule definition](#)


Note: You must [schedule custom reports](#) from the **Report > Custom Reports** tab.

Export Report Data

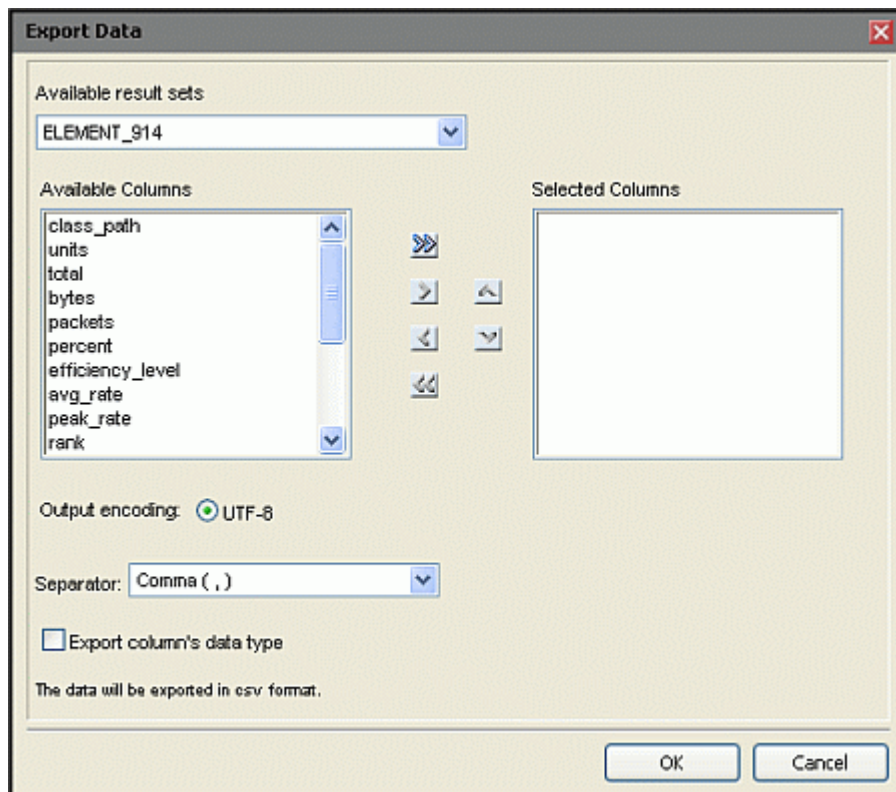
When you view a report in the IntelligenceCenter HTML report viewer, you can export the report data to a comma-separated values (CSV) file. You can then use the data to create your own documents in other applications that support CSV files, such as Microsoft Excel.



To export report data:

1. View a report in the HTML viewer, either by running an HTML version of the report using the [Run and View](#) option or by clicking the HTML icon in the **Action** column of the [Archived Reports](#) list.
2. Click the Export data button in the report viewer.



The *Export Data* dialog box is displayed.



3. Select the result set you want to export from the **Available result sets** drop-down list. In an IC report, each element that is displayed on the report—such as the title, subtitle, reporting time range, pie charts, tables, and graphs—has a corresponding result set entry. However, only the actual data elements such as tables and pie charts actually contain data that you would want to export. When you select a result set, the corresponding columns are displayed in the **Available Columns** box. If you select a result set that does not have any data associated with it, you will see a single item in the **Available Columns** box, such as Title or Subtitle. It will take some trial and error to select the result set that contains the actual data, but when you do you will see column names that correspond to actual data (such as bytes, avg_rate, and peak_rate) as shown in the dialog above.
4. Select the data columns that you want to export.
 - To select all columns, click the Add all  button.
 - To select columns individually, click the column you want (or Ctrl+Click or Shift+Click to select multiple columns) and then click the Add  button.
5. Click **OK**.
6. When prompted, save the CSV file to your computer. You can now open the file in an external application.

Manage Scheduled Tasks

In IntelligenceCenter, scheduled tasks represent the static reports that you have scheduled (or reports that are scheduled run by default). You can view the scheduled reports or remove a scheduled report if it is no longer needed or if you want to change the configuration.

To view and/or delete reports that were previously scheduled:

1. Select **Manage > Scheduled Tasks**. A list of all the scheduled reports is displayed (including [custom reports](#)). Note that this list includes reports that you have scheduled as well as reports that were scheduled to run by default.

Scheduled Tasks				
	Top N Applications_System Weekly <i>Report process for the weekly TopNApplications</i>	Waiting	Last Run: Nov 16, 2008 12:35 AM	Next Run: Nov 23, 2008 12:35 AM
	Top N Applications_System Monthly <i>Report process for the monthly TopNApplications</i>	Waiting		Next Run: Dec 1, 2008 1:30 AM
	Top N Talkers_System Yearly <i>Report process for the yearly TopNTalkers</i>	Waiting		Next Run: Jan 1, 2009 2:00 AM
	Top N Listeners_System Weekly <i>Report process for the weekly TopNListeners</i>	Waiting	Last Run: Nov 16, 2008 12:15 AM	Next Run: Nov 23, 2008 12:15 AM
	Top N Applications_System Daily <i>Report process for the daily TopNApplications</i>	Waiting	Last Run: Nov 20, 2008 12:40 AM	Next Run: Nov 21, 2008 12:40 AM
	Top N Traffic Classes Summary_System Daily <i>Report process for the daily TopNTrafficClasses</i>	Waiting	Last Run: Nov 20, 2008 12:50 AM	Next Run: Nov 21, 2008 12:50 AM
	Top N Traffic Classes Summary_System Yearly <i>Report process for the yearly TopNTrafficClasses</i>	Waiting		Next Run: Jan 1, 2009 2:40 AM

2. To view details about a specific scheduled report, click the right-arrow icon next to the report entry. The report entry expands to show detailed information about the report and the schedule on which it is configured to run.

	Top N Applications_System Weekly Report process for the weekly TopNApplications	Waiting	Last Run: Nov 16, 2008 12:35 AM	Next Run: Nov 23, 2008 12:35 AM
Schedule Details		Parameter Values		
Start	Nov 9, 2008 12:00 AM	Max Count	10	
End	No end date	Unit of Measure	bytes	
Recurrence:	Every week on Sunday	Sort Column	bytes	
		Sort Direction	desc	
		Show Guaranteed Rate	0	
		Failures		
Deactivate Delete				

3. If you want to stop a scheduled report from running, do one of the following:
 - To permanently remove the scheduled report, click the **Delete** link in the expanded report entry section.
 - To temporarily stop a scheduled report from running, click the **Deactivate** link. When you want to resume running the scheduled report, you can go back in and click the **Activate** link.

Note: If you want to change the configuration of a scheduled report, you must delete it and then [schedule the report](#) again.
4. To hide the schedule details, click the down-arrow icon icon next to the report entry.

Manage Report Schedules

Add a Schedule Definition

IntelligenceCenter provides some basic schedule definitions that you can use to schedule your reports. However, you may want to create additional schedule definitions that match the reporting schedules that are required within your organization. For example, you may want to add quarterly or semi-annual schedule definitions. To add a new schedule definition:

1. Select **Configure > Schedules**.
2. Click **Add** at the bottom of the *Schedule Definitions* pane. The *Create Schedule* dialog box opens.
3. Enter a schedule **Name**.
4. Enter a schedule **Description** (optional).
5. Specify the start and stop dates for the schedule by entering values in the **Range** section as follows:
 - To set the date and time that you want reports using this schedule to begin generating, click **Select** in the **Start** field. The *Date and Time* dialog box is displayed. Select the start date and time and then click **OK**.

Set when you want reports using this schedule to stop generating; reports that use this schedule will not run after the end that you specify. By default, the system defined schedules do not have an end date; that is, they will continue to run on the defined schedule indefinitely. If you want the reports using this schedule to run only a certain number of times, select **End after** and specify the number in the corresponding field. If you want the reports using this schedule to run only up to a specific date, select **End by** and then click **Select** in that field. The *Date and Time* window is displayed. Select the end date and time and then click **OK**.

6. Specify the specific pattern on which reports using this schedule will recur by setting the values in the **Recurrence** section as follows:
 - If you do not want the report to recur, select **No recurrence** (the default).
 - If you want the report to recur, select the frequency on which the recurrence pattern is based: **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**.
 - Specify the pattern. The specific parameters you can define depend on the frequency you selected. For example, if you selected **Monthly** as the frequency value, you can specify the number of months after which to run the report and the specific day of the month on which to run it. The following table describes the parameters associated with each frequency:

Frequency	Associated Parameters
Hourly	Allows you to define a recurrence pattern based on hour-long increments. For example, you can define a schedule that runs every 3 hours, 6 hours, or 9 hours (up to 24). By default, reports will run every hour.
Daily	Allows you to define a recurrence pattern based on day-long increments. For example, you could define the schedule to run every day, every other day, every 100 days (up to 365). By default, reports will run every day.
Weekly	Allows you to define a recurrence pattern based on week-long increments, up to 52 weeks. For example, you can define the schedule to run weekly, bi-weekly (every 2 weeks), or quarterly (every 13 weeks). By default, reports will run every week. You can also specify the specific day of the week on which you want the report to run. Selecting multiple days of the week indicates that the report will be run multiple times on the weekly pattern you've set. For example, if you set the report to run Every 2 weeks on Monday and Friday, the report will run twice every other week. If you want the report to run every weekday (Monday through Friday), click the Weekdays link.
Monthly	Allows you to define a recurrence pattern based on month-long increments, up to 12. In addition, you can specify the specific day of the month that you want the report to run. For example, to set up a semi-annual schedule that would run the 15th day of the month, you would set up the Monthly pattern to run every 6 months on day 15. By default, the reports run using this schedule will run the first of every month.
Yearly	Allows you to define a schedule that you can use to run reports once a year. You can specify the specific month and day that reports using this schedule will run.

7. When you are done configuring the new schedule definition, click **Save**.

Define Custom Schedule

To define a custom schedule:

1. Enter a schedule **Name**.
2. Enter a schedule **Description** (optional).
3. Specify the start and stop dates for the schedule by entering values in the **Range** section as follows:
 - To set the date and time that you want reports using this schedule to begin generating, click **Select** in the **Start** field. The *Date and Time* dialog box is displayed. Select the start date and time and the reporting time zone and then click **OK**.



- Set when you want reports using this schedule to stop generating; reports that use this schedule will not run after the end that you specify. By default, the system defined schedules do not have an end date; that is, they will continue to run on the defined schedule indefinitely. If you want the reports using this schedule to run only a certain number of times, select **End after** and specify the number in the corresponding field. If you want the reports using this schedule to run only up to a specific date, select **End by** and then click **Select** in that field. The *Date and Time* dialog box is displayed (see above). Select the end date and time and the reporting time zone and then click **OK**.
4. Specify the specific pattern on which reports using this schedule will recur by setting the values in the **Recurrence** section as follows:
 - If you do not want the report to recur, select **No recurrence**.
 - If you want the report to recur, select the frequency on which the pattern is based: **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**.
 - Specify the pattern. The specific parameters you can define depend on the frequency you selected. For example, if you selected **Monthly** as the frequency value, you can specify the number of months after which to run the report and the specific day of the month on which to run it. The following table describes the parameters associated with each frequency:

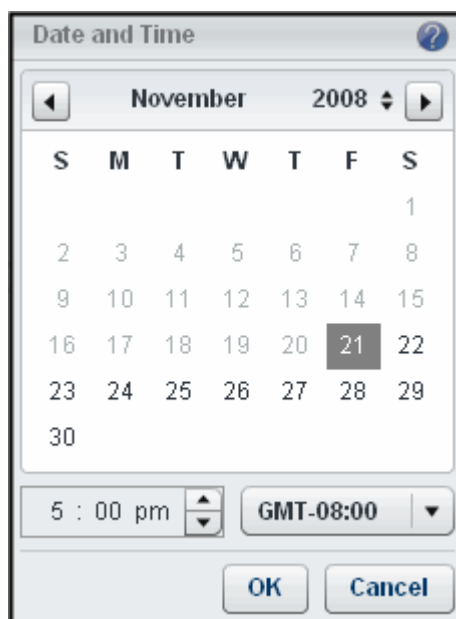
Frequency	Associated Parameters
Hourly	Allows you to define a recurrence pattern based on hour-long increments. For example, you can define a schedule that runs every 3 hours, 6 hours, or 9 hours (up to 24). By default, reports will run every hour.
Daily	Allows you to define a recurrence pattern based on day-long increments. For example, you could define the schedule to run every day, every other day, every 100 days (up to 365). By default, reports will run every day.
Weekly	Allows you to define a recurrence pattern based on week-long increments, up to 52 weeks. For example, you can define the schedule to run weekly, bi-weekly (every 2 weeks), or quarterly (every 13 weeks). By default, reports will run every week. You can also specify the specific day of the week on which you want the report to run. Selecting multiple days of the week indicates that the report will be run multiple times on the weekly pattern you've set. For example, if you set the report to run Every 2 weeks on Monday and Friday, the report will run twice every other week. If you want the report to run every weekday (Monday through Friday), click the Weekdays link.
Monthly	Allows you to define a recurrence pattern based on month-long increments, up to 12. In addition, you can specify the specific day of the month that you want the report to run. For example, to set up a semi-annual schedule that would run the 15th day of the month, you would set up the Monthly pattern to run every 6 months on day 15. By default, the reports run using this schedule will run the first of every month.
Yearly	Allows you to define a schedule that you can use to run reports once a year. You can specify the specific month and day that reports using this schedule will run.

- When you are done configuring the new schedule definition, click **Save**.

Modify a Schedule Definition

To modify a schedule definition:

1. Select **Configure > Schedules**.
2. Select the schedule definition that you want to modify from the *Schedule Definitions* pane. The schedule details are displayed in the right-hand pane.
3. If you want to modify the start and stop dates for the schedule, modify the values in the **Range** section as follows:
 - To set the date and time that you want reports using this schedule to begin generating, click **Select** in the **Start** field. The *Date and Time* dialog box is displayed. Select the start date and time and the reporting time zone and then click **OK**.



- Set when you want reports using this schedule to stop generating; reports that use this schedule will not run after the end that you specify. By default, the system defined schedules do not have an end date; that is, they will continue to run on the defined schedule indefinitely. If you want the reports using this schedule to run only a certain number of times, select **End after** and specify the number in the corresponding field. If you want the reports using this schedule to run only up to a specific date, select **End by** and then click **Select** in that field. The *Date and Time* dialog box is displayed (see above). Select the end date and time and the reporting time zone and then click **OK**.
4. If you want to modify the specific pattern on which reports using this schedule will recur, modify the values in the **Recurrence** section as follows:
 - If you do not want the report to recur, select **No recurrence**.
 - If you want the report to recur, select the frequency on which the pattern is based: **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**.
 - Specify the pattern . The specific parameters you can define depend on the frequency you selected. For example, if you selected **Monthly** as the frequency value, you can specify the number of months after which to run the report and the

specific day of the month on which to run it. The following table describes the parameters associated with each frequency:

Frequency	Associated Parameters
Hourly	Allows you to define a recurrence pattern based on hour-long increments. For example, you can define a schedule that runs every 3 hours, 6 hours, or 9 hours (up to 24). By default, reports will run every hour.
Daily	Allows you to define a recurrence pattern based on day-long increments. For example, you could define the schedule to run every day, every other day, every 100 days (up to 365). By default, reports will run every day.
Weekly	Allows you to define a recurrence pattern based on week-long increments, up to 52 weeks. For example, you can define the schedule to run weekly, bi-weekly (every 2 weeks), or quarterly (every 13 weeks). By default, reports will run every week. You can also specify the specific day of the week on which you want the report to run. Selecting multiple days of the week indicates that the report will be run multiple times on the weekly pattern you've set. For example, if you set the report to run Every 2 weeks on Monday and Friday, the report will run twice every other week. If you want the report to run every weekday (Monday through Friday), click the Weekdays link.
Monthly	Allows you to define a recurrence pattern based on month-long increments, up to 12. In addition, you can specify the specific day of the month that you want the report to run. For example, to set up a semi-annual schedule that would run the 15th day of the month, you would set up the Monthly pattern to run every 6 months on day 15. By default, the reports run using this schedule will run the first of every month.
Yearly	Allows you to define a schedule that you can use to run reports once a year. You can specify the specific month and day that reports using this schedule will run.

- When you are done modifying the schedule definition, click **Save**.

Create Consolidated Reports

The consolidated reporting feature allows you to configure report options for a sequence of reports—both top-level reports as well as drill-down reports—and schedule them to run together and be output as a single report. This feature does not allow you to change the format of the existing reports, rather it allows you to configure a series of standard reports and run them as a single report.

This feature is useful for many reasons. For example, perhaps you want to monitor the Top Applications in all of your Network Groups. Instead of running a report for each Network Group separately, you can create a custom report that includes an instance of the Top Applications Report for each of your Network Groups. You can then schedule this consolidated report to run and you will now be able to see the comparison in a single report rather than having to look in several different reports.

This feature also allows you to run reports at any level of the report hierarchy. For example, suppose you want to be able to monitor the TCP health of your Inbound/HTTP and Outbound/HTTP traffic daily. You could create a consolidated report that contains just the level two reports for these traffic classes. Then instead of generating the two separate reports and drilling-down to the information you are interested, you could view it in a single report.

In addition, this feature allows you to concatenate information from separate network groups (and hence, different DataCollectors) into a single report.

To create a consolidated report, you must:

- [Define the contents the consolidated report](#). You do this by configuring and running the reports you want to add to it. This process is called *recording*.
- [Schedule the consolidated report to run](#).

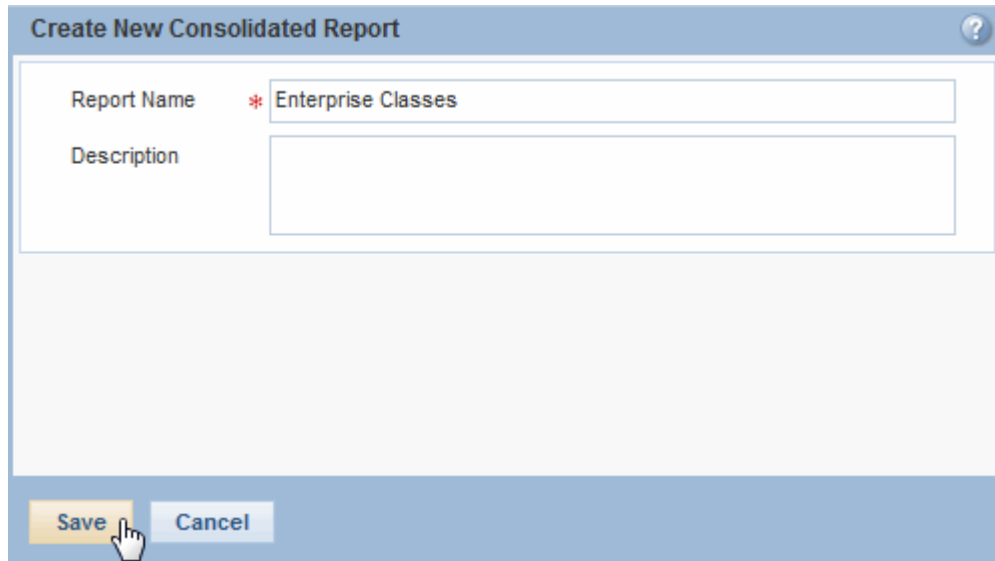
After they run, you can [view consolidated reports](#) from the Archived Reports pane just like you would view any other report.

Define a Consolidated Report

When you define a consolidated report, you must configure and run the reports you want to add to it. Note that you do not have to add all levels of a report to your consolidated report. Also, you can add multiple instances of the same level of a report, for instance if you drill-down on different applications, classes, or devices. In addition, consolidated reports are not limited to a single network group (and hence a single DataCollector). You can run the same report in different network groups or for different devices or views. This process of running reports and adding them to a consolidated report is called *recording*.

To define a consolidated report:

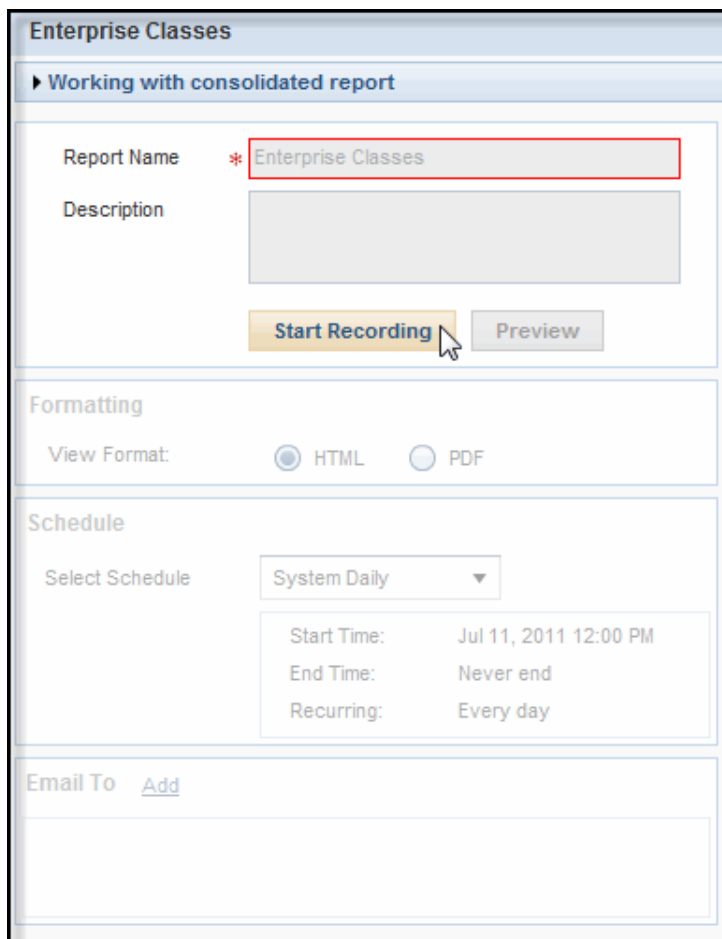
1. Select **Report > Consolidated Reports**.
2. Click **Add**. The Create New Consolidated Report dialog displays.



The dialog box titled "Create New Consolidated Report" has a blue header bar with a question mark icon on the right. It contains two input fields: "Report Name" with a red asterisk icon and the text "Enterprise Classes", and "Description" which is empty. At the bottom, there are two buttons: "Save" (highlighted with a mouse cursor) and "Cancel".

3. Enter a **Report Name** and optionally a **Description** and then click **Save**. A configuration area for the new custom report displays in the middle pane.

Note: The report name you specify will be the title that is displayed on the consolidated report.



The configuration pane titled "Enterprise Classes" has a blue header bar. Below the header is a section titled "Working with consolidated report" with a blue arrow icon. It contains the same "Report Name" and "Description" fields as the dialog box, with "Report Name" containing "Enterprise Classes". Below these fields are two buttons: "Start Recording" (highlighted with a mouse cursor) and "Preview".


Below the "Working with consolidated report" section is a "Formatting" section with a "View Format:" label and two radio buttons: "HTML" (selected) and "PDF".

Below the "Formatting" section is a "Schedule" section with a "Select Schedule" label and a dropdown menu showing "System Daily". Below the dropdown are three rows of information: "Start Time: Jul 11, 2011 12:00 PM", "End Time: Never end", and "Recurring: Every day".

At the bottom is an "Email To" section with an "Add" link and an empty text box.

4. Click **Start Recording**. After you click, the button name changes to **Stop Recording**.
5. Select **Run Reports** and select the first report you want to include in your custom report.
6. Configure the report as you want it to generate in your custom report and then click **Run**. When configuring a report for inclusion in a consolidated report, you must select **Run and View** from the **Run Options** field and **HTML** from the **View Format** field. You can configure all other fields as desired.

Note: When you are in the recording mode, the time ranges that are available on the Report Configuration screen are limited to **Year**, **Month**, **Week**, **Day**, and **Custom**.



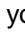



7. When the report displays, you will see an **Add to Consolidated Report** button in the upper left corner of the report.
 - To add the current report level to the report, click **Add to Consolidated Report**.
 - To skip this report level, drill-down to a lower-level report and then click **Add to Consolidated Report**.
 - You can add as many reports (drill-down or top-level reports) as you want to the custom report. Each time you click **Add to Consolidated Report**, the following confirmation displays:



8. When you are done adding reports, go back to the browser window or tab where IntelligenceCenter is running and select **Report > Consolidated Reports**. All of the reports you have added to the custom report so far are displayed in the Member Reports pane.



9. Review the list of reports and make the following modifications if necessary. Note that you will not be able to make any changes after you stop recording.:
 - If you want to add more reports, repeat steps 5 through 7.
 - To delete one of the reports, click the  icon next to the report name.
 - To rearrange the report order, select a member report and then click  to move it up or  to move it down. You will not be able to rearrange them after you stop recording.
 - To view configuration details for one of the reports, click the  icon in the Member Reports pane. If you want to change the report configuration, you must delete it and then re-add it. To customize the description of the report that displays in the Member Reports pane, click **Edit** and then enter a description that uniquely identifies the report, such as Level 2 Inbound HTTP Class Utilization. When you are done modifying the description, click **Save**.

Report Details

Class Utilization

Parameter Values

Traffic Class	/Inbound/Default
Consider Business Hours	false
Sort Column	bytes
Time Span	7 Day(s)
Sort Direction	desc
Max Count	10
Network Group or Device	Network
Report Description	

[Edit](#)

[Close](#)

- If you are happy with the list of reports, continue to the next step.
10. If you want to see what the consolidated report will look like, click **Preview**. IC will run all of the reports you have added so far and display them in the consolidated format.
 11. When you have finished adding reports and arranging them in the desired order, click **Stop Recording**. The Formatting, Schedule, and Email To sections of the screen become active when you stop recording. You can now [schedule the new custom report](#).

Report Name * Enterprise Classes

Description

[Stop Recording](#) [Preview](#)

Formatting

View Format: ☒ HTML ☐ PDF

Schedule

Select Schedule: System Daily ▼

Start Time: Jul 11, 2011 12:00 PM



End Time: Never end

Recurring: Every day

Email To [Add](#)

Schedule a Consolidated Report

After you create your consolidated report, you must schedule it to run as follows:

1. Select **Report > Consolidated Reports**.
2. In the Consolidated Reports pane, select the consolidated report you want to schedule.
3. Select **HTML** for **View Format**.
4. In the **Select Schedule** field, select the schedule you want to use to run this report, or select **Create Schedule** to [define a custom schedule](#) for this report for this report.
5. To email the report after it runs, click the **Add** link in the Email To field and then define recipients as follows:
 - If you want to email the report to an existing IntelligenceCenter user, select the user from the **Available** column and click the  button to move the user to the **Selected** column. To select multiple, consecutive users Shift+Click. To add multiple, non-consecutive users, Ctrl+Click.
 - If you want to email the report to a recipient who does not have an IntelligenceCenter [user profile](#), enter the user's **Email** address and optionally the user's **Name** and then click the add  button. The recipient is added to the **Selected** list. Repeat this step for each user you want to add. You can add any email address that is recognized by your SMTP server, including group alias addresses. Note, however, that any external users you add are not saved to the IntelligenceCenter database and you will have to readd them each time you run a report that you want to send to that person. Therefore, if you find that you are sending reports to the same people frequently, you should [create a user profile](#) for the person.
 - When you are done defining the recipients, click **OK**. The recipients you added are displayed on in the *Email To* section of the *Report Configuration* pane. Note that you can delete a recipient by clicking the **Remove** link next to the recipient's user profile name or email address.

6. Click **Schedule**. An entry for the report will be added to the [Scheduled Tasks](#) list.

The screenshot shows the 'Enterprise Classes' report configuration interface. The left pane, 'Consolidated Reports', lists 'cr' and 'Enterprise Classes'. The right pane, 'Enterprise Classes', contains the following sections:

- Working with consolidated report:** Includes 'Report Name' (set to 'Enterprise Classes'), a 'Description' text area, and 'Start Recording' and 'Preview' buttons.
- Formatting:** Includes a 'View Format' section with radio buttons for 'HTML' (selected) and 'PDF'.
- Schedule:** Includes a 'Select Schedule' dropdown set to 'System Daily', and fields for 'Start Time' (Jul 11, 2011 12:00 PM), 'End Time' (Never end), and 'Recurring' (Every day).
- Email To:** Includes an 'Add' link and a text area.

At the bottom of the right pane, there are buttons for 'Add', 'Remove', 'Refresh', 'Schedule' (highlighted with a mouse cursor), and 'Reset'.

7. After the report runs, an entry will appear on the Archived Reports pane. You can [view reports](#) in the Archived Reports pane at any time. Consolidated reports display with the consolidated report name and the time. Hover over the report name to view details.

Manage Sites

In IntelligenceCenter, a *site* is a named reference to a subnet class on a PacketShaper traffic tree. Sites are frequently associated with locations: branch offices, departments, and so forth. Traffic trees that are organized in this fashion are commonly referred to as *location-based* or *site-based* trees. By enabling discovery on these subnet classes, you can identify what applications are running at each site and assign appropriate policies. For more information about creating site-based traffic trees, see [Create a Location-Based Traffic Tree with Per-Location Applications](#) in PacketGuide.

When your PacketShaper appliances are configured with site-based traffic trees, you can use IC site reports to display data associated with each of these pre-defined sites. For example, a PacketShaper might classify the network traffic by subnet, with each location (subnet) having its own traffic class (Cupertino, Sunnyvale, San_Jose, and so forth). In IC, you can create reports that show statistics for all sites on your network, as well as create reports for a specific site.

Before you can [run site-based reports](#), you must first associate your subnet classes with an IC site name. You can perform the following site management tasks:

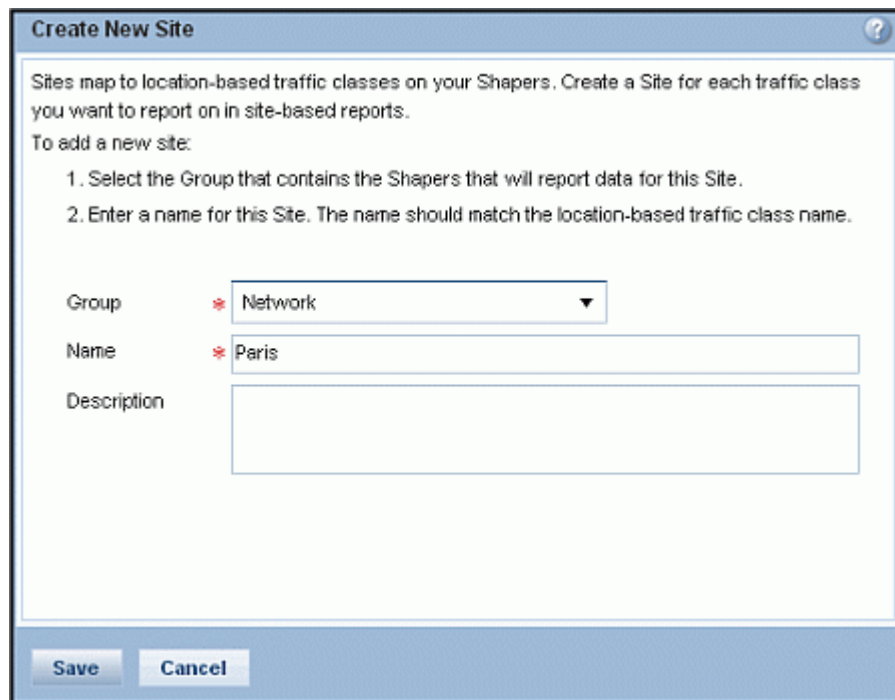
- [Add a Site](#)
- [Remove a Site](#)

Add a Site

In order to [run site-based reports](#), you must first create IntelligenceCenter site names that are associated with each subnet class.

To add a site to IC:

1. Select **Configure > Network > Sites**.
2. Click **Add** at the bottom of the *Sites* pane. The *Create New Site* dialog box is displayed.



Create New Site

Sites map to location-based traffic classes on your Shapers. Create a Site for each traffic class you want to report on in site-based reports.

To add a new site:

1. Select the Group that contains the Shapers that will report data for this Site.
2. Enter a name for this Site. The name should match the location-based traffic class name.

Group * Network ▼

Name * Paris

Description

Save Cancel

3. Select the network **Group** that contains the PacketShaper appliances that report data for the site.
4. Enter the **Name** of the site. The site name must exactly match the traffic class name used on the PacketShaper.
5. Enter a **Description** for the site. It's good practice to include the name of the PacketShaper on which the site class is located in the description.
6. Click **Save**. The new site name will be listed alphabetically in the *Sites* pane.

Remove a Site

If you no longer want to report on a site or if the associated subnet class no longer exists on the PacketShaper, you can remove the site from IntelligenceCenter.

To remove a site from IC:

1. Select **Configure > Network > Sites**.
2. Select the site you want to remove and click **Remove**.
3. When prompted to confirm the site removal, click **Yes**. The site is removed from the **Sites** tab.

After a site is removed, the name will no longer be available for selection when running site-based reports.

Manage Alerting

Whenever there is an IC or DC event that may require administrator attention, it can trigger an alert notification to be logged to the IC alert viewer, [sent via an email](#), [SNMP trap](#), [Syslog server](#), or a combination of any of these methods. IC includes a set of predefined alert type definitions that define the events that trigger alert notification. By default, all of the predefined alert types are enabled for notification and do not require additional configuration. However, you can customize specific alert type definitions, or you can configure alerts globally by severity type. You can also [enable or disable the alerting system](#).

You can perform the following alerting tasks:

- [Enable or disable alerting](#)
- [View alerts](#)
- [Configure alerts by severity type](#)
- [Configure individual alert type definitions](#)
- [Enable syslog](#)
- [Enable SNMP traps](#)

Alert Type Configuration Details

An alert is a notification that a specific event has occurred on the IC or DC system. The definition of the events that trigger alerts are called alert types. Each alert type definition includes a name and optionally a description. In addition, it includes the following configuration information:

- **Alert Title**—The name that appears on the **View Alerts** screen when an instance of the event defined by alert type occurs. Note that by default, the Alert Title is the same as the Alert Type; changing the Alert Title, however, does not change the name of the Alert Type.
- **Alert State**—Whether the alert type is **On** or **Off**. By default, all alert types are **On** and will therefore will trigger alert notification when the corresponding event occurs on the system.
- **Severity**—A classification that indicates the seriousness of the event that caused the alert. The severity level can be **Informative**, **Warning**, **Minor**, **Major**, **Critical**. All alert types have a default severity, but you can change the value depending on which events are most critical within your organization. You can also filter which alerts are displayed on the View Alerts screen based on severity level and you can configure other settings based on severity level.
- **Description**—Describes the event that triggers the alert.
- **Auto deletes after**—Specifies the number of days after which the alert will automatically be deleted from the database. By default, all alerts are configured to auto delete after 30 days.
- **Coalesces occurrences within**—Sometimes, an event that causes an alert notification will occur multiple times within a short period. To prevent, a flood of duplicate alert instances, only the first event triggers an alert notification for the duration of the coalesce interval. For example, if the coalesce interval is set to coalesce instances within ten minutes, and the event that triggers the event reoccurs within five times within that ten minute interval, only one alert notification will be generated. However, the details of the alert notification will show the number of times the event occurred (the Count) as well as the date and time of the first occurrence and the last occurrence. If the same event that

triggered the alarm occurs again fifteen minutes after the initial event, a second alert notification will be generated because it is outside of the coalesce interval. You can set coalesce intervals by severity level or you can set them individually for a given alert type.

- **Acknowledgement method**—Specifies who has to acknowledge an alert for it to be considered acknowledged. When this value is set to **Any**, the alert will be considered acknowledged as soon as a single user in its configured subscription group acknowledges it. When this value is set to **All**, the alert will not be considered acknowledged until all users/groups in its configured subscription group have acknowledged it.
- **My Email Subscription**—Enables email notification of alerts for the logged in user. Note that you must have configured the mail server before you can use email notification.

View Alerts

Whenever there is an IC or DC event that corresponds to an alert type that is enabled on the system, an alert notification is generated. Every alert notification is displayed on the View Alerts screen. Depending on the [alert type configuration](#), the event may also trigger notification by email, syslog, or SNMP trap.

To view alert notifications from within IC:

1. Select **Alert > View Alerts**. The View Alerts screen is displayed. A summary of the alerts is displayed at the top of the screen.

Alert Title	Source	Severity	Count	First Occurrence	Last Occurrence	Acknowledge	Save	Delete
DataCollector exceeded memory threshold	DataCollector1	Critical	4	Jan 10, 2010 12:05 PM	Jan 10, 2010 12:12 PM	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Each alert shows the following information:

- The alert title
 - The source, which is the device on which the event occurred
 - Severity of the alert (by default, only alerts with a severity level of Major or Critical are displayed, but this is configurable)
 - The count, which is the number of times the event has occurred
 - The date and time of the first and last occurrence of the event that triggered the alert
2. To view more details about a specific alert instance, click the corresponding link in the **Alert Title** field. The *Alert Instance Details* dialog box opens. This dialog box shows a description of the alert as well as any comments that users have entered about the alert instance. To enter a comment, click the **Add Comment** button.

Alert Instance Details ? X

Information

Source: DataCollector1
 Severity: Critical
 First Occurrence: Jan 10, 2010 12:05 PM
 Last Occurrence: Jan 10, 2010 12:12 PM
 Description: DataCollector : DataCollector1 exceeded configured memory threshold

History

Total Occurrences: 4
 Source: DataCollector1
 Acknowledged: false

Acknowledge Alert: ☐

No comments

Add Comment

Cancel **Save**

3. If you want to change the number of alerts that are displayed (20 are displayed by default), select a value from the **Alerts per Page** field. If there are more alerts that can be displayed on a single page, you can scroll through them using the arrow buttons or enter a page number in the text box.

4. Specify an action to take on the alert, if any:
 - To acknowledge the alert, check the **Acknowledge** checkbox.
 - To save the alert, check the **Save** checkbox. Saved alerts will remain on the alert list until you delete them.
 - To delete an alert, check the **Delete** checkbox.
5. If you want to filter the list of alerts that is displayed or to search for a specific alert, expand the **Filter Alert** section of the screen by clicking the right arrow ► icon:

▼ Filter Alerts

Source	Severity	Acknowledgement	Saved	Time Range
<input checked="" type="checkbox"/> DC	<input type="checkbox"/> Informative	<input type="checkbox"/> Acknowledged	<input type="checkbox"/> Saved	This Week ▼
<input checked="" type="checkbox"/> IC	<input type="checkbox"/> Warning	<input checked="" type="checkbox"/> UnAcknowledged	<input checked="" type="checkbox"/> Unsaved	
	<input type="checkbox"/> Minor			
	<input checked="" type="checkbox"/> Major			
	<input checked="" type="checkbox"/> Critical			

Reset **Apply Filters**

- Check the checkboxes that correspond to the alerts you want to display. You can filter by source (IC or DC), severity, acknowledgement status, or saved/unsaved.

Note that whether or not an alert is considered to be acknowledged depends on how the alert type is configured. In some configurations, the alert must be acknowledged by at least one user to be considered acknowledged; in other configurations, it must be acknowledged by all users in the notification group to be considered acknowledged.

- Select the time range of alerts you want to display by selecting a value from the **Time Range** field drop-down list. You can choose to display alerts from **Today**, **This Week** (the default), **Last Week**, **This Month**, **Last Month**, **This Quarter**, or **Last Quarter**.
- Click **Apply Filters**. The list of displayed alerts is updated to display only those alerts that match the criteria you entered. To set the filter back to the default values, click **Reset**.

Configure Global Alert Settings

An alert is a notification that a specific event has occurred on the IC or DC system. The definition of the events that trigger alerts are called alert types. Although you cannot add new alert type definitions to IC, you can modify the configuration of the existing ones. In addition, you can disable specific alert types if they do not pertain to your organization or you do not care about being notified of the specific event.

To configure global alert type settings:

1. Select **Alert > Configure Alerting**. The **Alert System** settings are displayed at the top of the screen. The bottom portion of the screen contains one of the alert configuration panes (**Alert Type Summary** or **Alert Types and Details**).
2. If the **Alert Type Summary Table** is not displayed on the left side of the bottom pane, click the **View Alert Type Summary** bar on the right side of the screen as shown below.

The screenshot shows the 'Device is offline' alert configuration interface. On the left, a tree view lists alert types under 'IC' and 'DC' categories. The 'Device is offline' alert type is selected. The main area displays configuration details for this alert type, including fields for Alert Title, Alert State, Alert Type, Severity, Description, Auto deletes after, Coalesces occurrences within, Acknowledgment Method, Notification Groups, and My Email Notification. A 'View Alert Type Summary' button is located on the right side of the screen.

The **Alert Type Summary** Table shows and allows you to edit the basic configuration settings all alert types.

3. To set the deletion interval for alerts based on alert severity, click **Deletion Intervals**. The *Set Global Deletion Interval* dialog box is displayed. Enter the number of days after which you want alerts to be deleted automatically in the **Set Deletion Intervals to** text box, select the severity levels to which you want this deletion interval to apply, and then click **Save**. Repeat this step to set the deletion interval for other severity levels.
4. To set the coalesce interval for alerts based on alert severity, click **Coalesce Intervals**. The *Set Global Coalesce Interval* dialog box is displayed. The coalesce interval specifies how often to trigger an alert when event that causes an alert notification occurs multiple times within a short period. Enter the number of minutes after which you want a duplicate event to trigger a new alert instance in the **Set Coalesce Intervals to** text box, select the severity levels to which you want this deletion interval to apply, and then click **Save**. Repeat this step to set the coalesce interval for other severity levels.
5. To enable or disable email notification for the logged in user based on severity level, click **My Email Notifications by Severity**. The *Set Global Notification by Severity* dialog box is displayed. Select the severity levels for which you want receive email notification and then click **Save**. This setting is saved with configuration settings for the user and therefore each user who wants to change the email notification setting must configure it individually.
6. If you want to modify [alert type configuration](#) settings for a specific alert type, select the corresponding row in the Alert Type Summary Table to make the fields editable.

Alert Type	Severity	Deletion Interval(Days)	Coalesce Interval(Mins.)	Acknowledgment	My Email Notification
▼ IC					
Device is offline	Major	30	10	All	<input checked="" type="checkbox"/>
External Authentication, Review account	Critical	30	15	All	On
IC Memory is low	Major	30	15	Any	On
IC lost contact with DC	Critical	30	15	All	On

Note: the summary table does not show all settings; to edit other configuration settings you must go to the [Alert Types and Details](#) screen.

Note: You cannot change the configuration settings on the **DataCollector license close to expiring** and **DataCollector license expired** alert types; you can only [modify notification group membership](#) on these alert types.

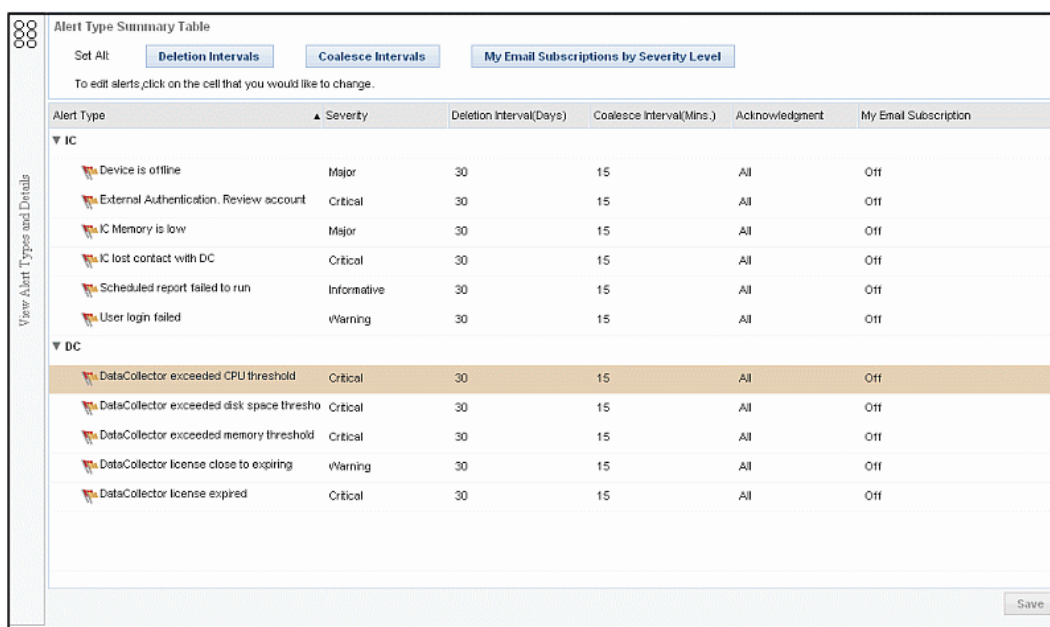
7. When you are done modifying the global alert type configuration settings, click **Save**.

Configure Alert Type Settings

An alert is a notification that a specific event has occurred on the IC or DC system. The definition of the events that trigger alerts are called alert types. Although you cannot add new alert type definitions to IC, you can modify the configuration of the existing ones. In addition, you can disable specific alert types if they do not pertain to your organization or you do not care about being notified of the specific event.

To configure alert type settings for a specific alert type:

1. Select **Alert > Configure Alerting**. The **Alert System** settings are displayed at the top of the screen. The bottom portion of the screen contains one of the alert configuration panes.
2. If the **Alert Types** pane is not displayed on the left side of the bottom pane, click the **View Alert Types and Details** bar on the left side of the screen as shown below.



Alert Type Summary Table

Set All: **Deletion Intervals** **Coalesce Intervals** **My Email Subscriptions by Severity Level**

To edit alerts, click on the cell that you would like to change.

Alert Type	▲ Severity	Deletion Interval(Days)	Coalesce Interval(Mins.)	Acknowledgment	My Email Subscription
▼ IC					
Device is offline	Major	30	15	All	Off
External Authentication. Review account	Critical	30	15	All	Off
IC Memory is low	Major	30	15	All	Off
IC lost contact with DC	Critical	30	15	All	Off
Scheduled report failed to run	Informative	30	15	All	Off
User login failed	Warning	30	15	All	Off
▼ DC					
DataCollector exceeded CPU threshold	Critical	30	15	All	Off
DataCollector exceeded disk space thresho	Critical	30	15	All	Off
DataCollector exceeded memory threshold	Critical	30	15	All	Off
DataCollector license close to expiring	Warning	30	15	All	Off
DataCollector license expired	Critical	30	15	All	Off

Save

3. Select the **Alert Type** you want to configure. Configuration details and a description of the selected alert type are displayed in the right pane.

Note: The **User login failed alert** is only triggered when a user with a valid user name fails to enter the correct password. It is not triggered when an invalid user name is entered and therefore is not intended to detect intrusion attempts on the system.

4. Modify the [alert type configuration](#) as desired.

Note: You cannot change the configuration settings on the **DataCollector license close to expiring** and **DataCollector license expired** alert types; you can only modify notification group membership on these alert types.

5. Set the **Acknowledgement Method**. If you set it to **All**, all users in the subscription group must acknowledge the alert in order for it to be considered acknowledged; if the field is set to **Any**, the alert instance will be considered acknowledged as soon as any member of the subscription group acknowledges it.
6. To enable or disable email notification for the logged in user, select the appropriate radio button in the **My Email Subscription** field. This setting is saved with the user's configuration settings and therefore each user that wants to change the email notification setting must configure it individually.
7. When you are done modifying the alert type configuration, click **Save**.

Set Global Coalesce Interval

To set the coalesce interval for alerts based on alert severity:

1. Click **Coalesce Intervals**. The *Set Global Coalesce Interval* dialog box is displayed. The coalesce interval specifies how often to trigger an alert when event that causes an alert notification occurs multiple times within a short period.
2. Enter the number of minutes after which you want a duplicate event to trigger a new alert instance in the **Set Coalesce Intervals to** text box.
3. Select the severity levels to which you want this deletion interval to apply.
4. Click **Save**.
5. Repeat Steps 1-4 to set the coalesce interval for other severity levels.

Set Global Deletion Interval

To set the deletion interval for alerts based on alert severity:

1. Click **Deletion Intervals**. The *Set Global Deletion Interval* dialog box is displayed.
2. Enter the number of days after which you want alerts to be deleted automatically in the **Set Deletion Intervals to** text box.
3. Select the severity levels to which you want this deletion interval to apply.
4. Click **Save**.
5. Repeat Steps 1-5 to set the deletion interval for other severity levels.

Set Global Notification by Severity

To enable or disable email notification for the logged in user based on severity level:

1. Click **My Email Notifications by Severity**. The *Set Global Notification by Severity* dialog box is displayed.
2. Select the severity levels for which you want receive email notification and then click **Save**. This setting is saved with configuration settings for the user and therefore each user who wants to change the email notification setting must configure it individually.

Send Alert Notifications to a Syslog Server

If the alerting system is enabled, you can configure it so that events that trigger alert notification result in syslog messages (provided that you have configured at least one [external syslog server](#)).

To send alert notifications to a syslog server:

1. Select **Alert > Configure Alerting**.
2. Check the checkbox in the **Enable Syslog for** field. This activates the drop-down list.



☒ Enable Syslog for: Informative ▼ or higher
[Configure Syslog](#)

Save

3. Select the alert severity level for which you want to send syslog messages from the drop-down list and then click **Save**. The IC server will now send a syslog message to the [configured syslog servers](#) whenever an alert with the selected severity level or higher is generated. Note that IC alert severities do not necessarily correlate to syslog severity levels.

Send SNMP Traps

If the alerting system is enabled, you can configure it so that events that trigger SNMP traps. Note that you must [configure at least one SNMP trap manager](#) before you can use this feature.

To enable SNMP traps:

1. Select **Alert > Configure Alerting**.
2. Check the checkbox in the **Enable SNMP trap for** field. This activates the drop-down list.



3. Select the alert severity level that you want to trigger SNMP traps from the drop-down list and then click **Save**. The IC server will now send an SNMP trap to the configured SNMP trap managers whenever an alert with the selected severity level or higher is generated.

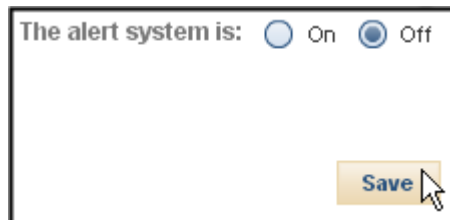
Enable or Disable Alerting

The IC alerting system is enabled by default. You can enable or disable alerting, which enables or disables the alerting system. If alerting is enabled, you can configure which events trigger alert notification either by [configuring alert types individually](#) or by [enabling or disabling alerts globally by severity level](#).

In addition, if alerting is enabled, you can also choose to have alert notifications sent to one or more [external syslog servers](#) or [SNMP trap managers](#).

To enable or disable alerting:

1. Select **Alert > Configure Alerting**.
2. Enable or disable alerting by selecting the **On** or **Off** radio button in the **The alert system is** field and then clicking **Save**.



Configure Portlets

IntelligenceCenter allows you to define sets of portlets that enable custom views of your network. The portlets you define are contained in a *portlet view*. By default, IntelligenceCenter contains a single portlet view called *My View*. However, you can [add and delete portlet views](#) as necessary to organize the portlets you define. Each view can contain multiple *portlets*, which are customizable network monitoring applications. The portlets and portlet views that you define are unique to your IntelligenceCenter [user profile](#), allowing you to create your own management dashboard that shows the aspects of the network that are important to you. You can even [add](#) multiple instances of the same portlet to a view, configuring each instance with different parameters so that you can monitor different slices of your network.

IntelligenceCenter provides seven portlets that you can use to monitor network applications and devices and troubleshoot application problems:

- The [Application Performance portlet](#) shows the health and utilization of up to 10 [applications](#) of interest.
- The [Class Utilization portlet](#) allows you to monitor and compare the average and peak bandwidth usage of up to five traffic classes.
- The [Network Efficiency portlet](#) allows you to monitor the efficiency of the TCP traffic classes on your network.
- The [Per Subnet FDR portlet](#) displays statistics for traffic between two sites (such as two subnets) or on a specific virtual LAN (VLAN).
- The [Per Server FDR portlet](#) tracks recent traffic flows to your inside servers and allows you to keep an eye on traffic that you consider to be suspect.
- The [Top N Children portlet](#) allows you to monitor the relative portions of bandwidth allocated to the ten most active children classes of a selected traffic class.
- The [VoIP Performance portlet](#) allows you to monitor the performance of the RTP- and RTCP-based VoIP calls running on your network.

When you first add a portlet to a portlet view, it may not contain any data or it may contain data that is not meaningful in your organization because it uses predefined configuration settings. In order to display meaningful information, you must [configure the portlet](#) with the specific information that you want it to display. Additionally, as DataCollector collects data, it rolls it up into time-based tables for reporting. If the table that corresponds to the time span you selected does not have data rolled up into it yet, the portlet will not display any data.

To run portlets, you must have a valid DataCollector [license](#) and you must have installed and [configured DataCollector](#) to collect data from your PacketShaper and/or other network devices.

Manage Portlets

IntelligenceCenter allows you to define sets of portlets that enable custom views of your network. This section describes how to manage IntelligenceCenter portlets. It includes the following topics:


- [Add and delete portlet views](#)
- [Add portlets](#)
- [Configure portlets](#)
- [View portlets](#)
- [Delete portlets](#)
- [Arrange portlets](#)

Manage Portlet Views

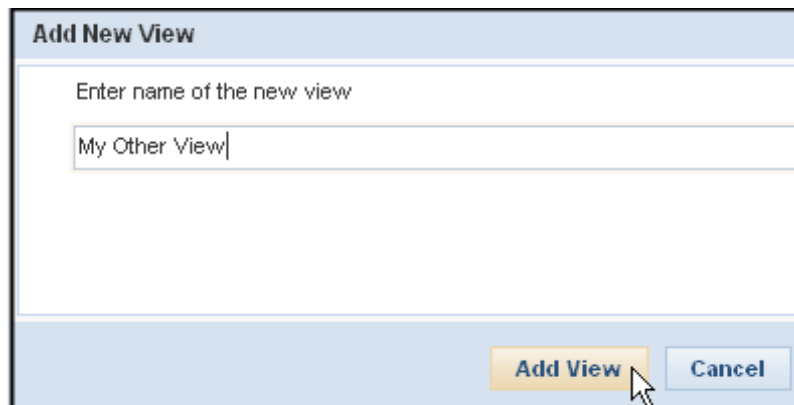
By default, the **Monitor** tab contains a single portlet view called *My View*. You can add and delete portlet views to suit your needs. For example, you could have one portlet view that allows you to monitor the health of your network. This view might contain several instances of the Application Performance portlet, each monitoring health on a different network sub-group or view. You could then have another view that displays statistical information for the various VLANs on your network.

Adding a Portlet View

To add a portlet view:

1. Click **Monitor**.
2. Click the  button and select one of the following options:
 - If you want to add the new portlet view as the first tab, select **Add New View to First**.
 - If you want to add the new portlet view as the last tab, select **Add New View to Last**.

The *Add New View* dialog box is displayed.

The image shows a dialog box titled "Add New View". It has a light blue header bar with the title. Below the header, there is a text input field with the placeholder text "Enter name of the new view". The field contains the text "My Other View". At the bottom of the dialog, there are two buttons: "Add View" (highlighted in orange) and "Cancel" (light blue). A mouse cursor is pointing at the "Add View" button.


3. Enter a name for the new view in the text field (up to 50 characters) and then click **Add View**.

4. The new view is displayed as a tab.



Deleting a Portlet View

To delete a portlet view:

1. Click **Monitor** in the IntelligenceCenter banner.
2. Select the tab for the view that you want to delete.
3. Click the  icon on the tab for the view you want to delete.



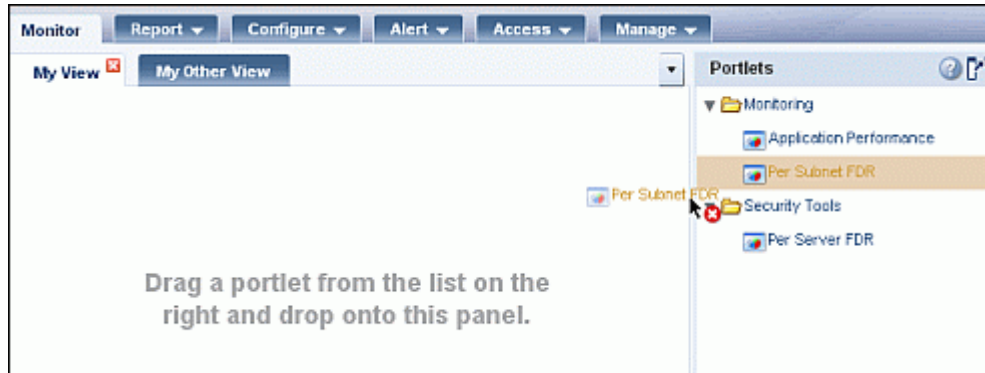
4. When prompted to confirm the deletion, click **Yes**.

Note: You cannot delete the last portlet view; there must always be at least one view.

Add a Portlet

To add a portlet:

1. Click **Monitor**.
2. Select the portlet view to which you want to add the portlet. For example, to add the portlet to the *My View* portlet view, click the **My View** tab.



3. Select the portlet you want to add in the *Portlets* pane and drag it to the selected portlet view. The portlet is displayed in the portlet view.

Note: You can add multiple instances of a portlet to the same portlet view. Newly added portlets will not display any data until you select a network group.

4. You can now:
 - [Configure the portlet](#)
 - [Arrange the portlet windows on the portlet view](#)

Configure a Portlet

You can customize the IntelligenceCenter portlets to show exactly the information that is important to you. The configuration changes you make to the portlets and portlet views are unique to your user profile. The procedure for configuring a portlet depends on the portlet:

- [Configure the Application Performance portlet](#)
- [Configure the Per Subnet FDR portlet](#)
- [Configure the Per Server FDR portlet](#)
- [Configure the Class Utilization Portlet](#)
- [Configure the Network Efficiency Portlet](#)
- [Configure the Top N Children Portlet](#)
- [Configure the VoIP Performance Portlet](#)


Configure the Application Performance Portlet

The [Application Performance](#) portlet can monitor the health and service level agreement (SLA) compliance of up to 10 TCP applications in a specified [network branch or view](#). The Application Performance portlet gives information about application performance based on three network quality parameters: loss, latency, and availability. IntelligenceCenter uses baseline values that you set for each of these parameters to calculate the performance of each application. In order for the Application Performance portlet to present meaningful compliance data, you must tune the parameters for each application you are monitoring according to how you expect the application to perform on your network.

When you configure the Application Performance portlet, you must consider three things:

- **What applications to monitor** — You can add and remove applications from the Application Performance portlet (or you can configure different instances of the Application Performance portlet to monitor different sets of applications). Keep in mind that in order to monitor the health of an application, IntelligenceCenter must know about the application. IntelligenceCenter includes applications that match most standard network traffic. However, if you have custom traffic classes or if you want to monitor non-standard applications, you will need to [create your own applications](#) or [modify the existing applications](#) to include your custom classes. In addition, only applications that use the TCP protocol will have valid availability, loss, and latency statistics; thus, only TCP applications will provide a valid performance calculation. UDP-based applications will not give meaningful results.
- **What baseline values to set for each application** — Baseline values for the loss, latency, and availability parameters should reflect the average values of these parameters on a normal day of operation on your network. It can also be useful to tune these parameters to correspond with service level agreements you have in place. You must set baseline values for loss, latency, and availability for each application you monitor in the Application Performance portlet. IntelligenceCenter will then compare the actual values of these parameters during the monitoring period to the baseline values you set to determine the percent difference, which it uses to calculate the performance compliance of the application.
- **How important each parameter is to the specific application** — The degree to which the loss, latency, and availability parameters is important in determining performance compliance depends on the application itself. For example, if you are monitoring a VoIP application, the latency parameter is much more important in determining the quality of the application than loss and should therefore be assigned a higher weight. You must assign weight values for loss, latency, and availability for each application you monitor in the Application Performance portlet.

To configure the Application Performance portlet:

1. Click **Monitor**.
2. Select the portlet view that contains the Application Performance portlet that you want to configure.
3. Click the  icon in the portlet's title bar and select **Configuration** from the pop-up menu. The **General** configuration tab appears.

Application Performance

General **Details**

Name: Application Performance

Select Time Range:

Today ▼

☐ Automatically update every 15 minutes

Apply Changes Cancel

4. If you want to customize the portlet name, enter a new **Name**.
5. Select the time range over which to monitor. By default, the portlet monitors application traffic collected today (rounded to the last full hour). However, you can select a different predefined time range from the drop-down menu or select **Other** to define your own time range. For best performance, you should use a predefined time range whenever possible.
6. If you want the portlet to refresh automatically, check the **Automatically update every xx minutes** checkbox and then set the frequency at which you want automatic updates to occur. The range is 15 minutes (the default) up to 1200 minutes.
7. Select the **Details** tab to configure the applications to monitor and their thresholds and to specify the network group to monitor. If you do not select a network group, the portlet will not display any data.

Application Performance

General | **Details**

Group: [Select...](#)

Performance Configurations [Add](#) [Collapse All](#) [Expand All](#)

▼ **CIFS** [Remove](#)

Application: [Select...](#)

Parameter	Baseline	Weight(%)
Latency:	30 ms	30%
Loss:	0.001	30%
Availability:	100 %	40%


[DNS](#) [Remove](#)
[DHCP](#) [Remove](#)
[Telnet](#) [Remove](#)
[HTTP](#) [Remove](#)

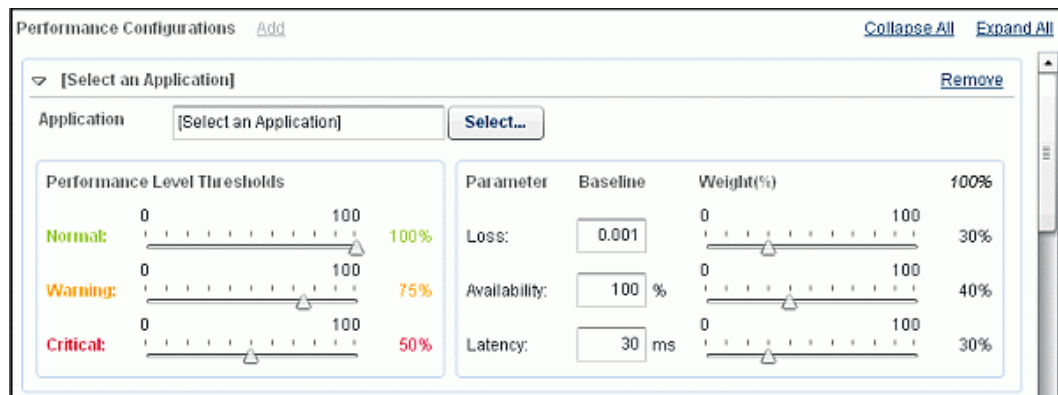
[Apply Changes](#) [Cancel](#)

8. Specify the network group or view for which the portlet should display data by clicking **Select** in the **Group** field. The *Select Group or View* dialog box is displayed. Select the network group, sub-group or device on the **Devices** tab or the logical network view on the **Views** tab to which to restrict this Application Performance portlet and then click **Accept**.
9. Customize the application definitions as desired. You can:
 - Tune parameter values on an existing application
 - [Add a new application to the Application Performance portlet](#)
 - Remove an application from the Application Performance portlet
10. When you are finished configuring the portlet, click **Apply Changes**.

Add a New Application Definition to the Application Performance Portlet

To add a new application definition to the Application Performance portlet:

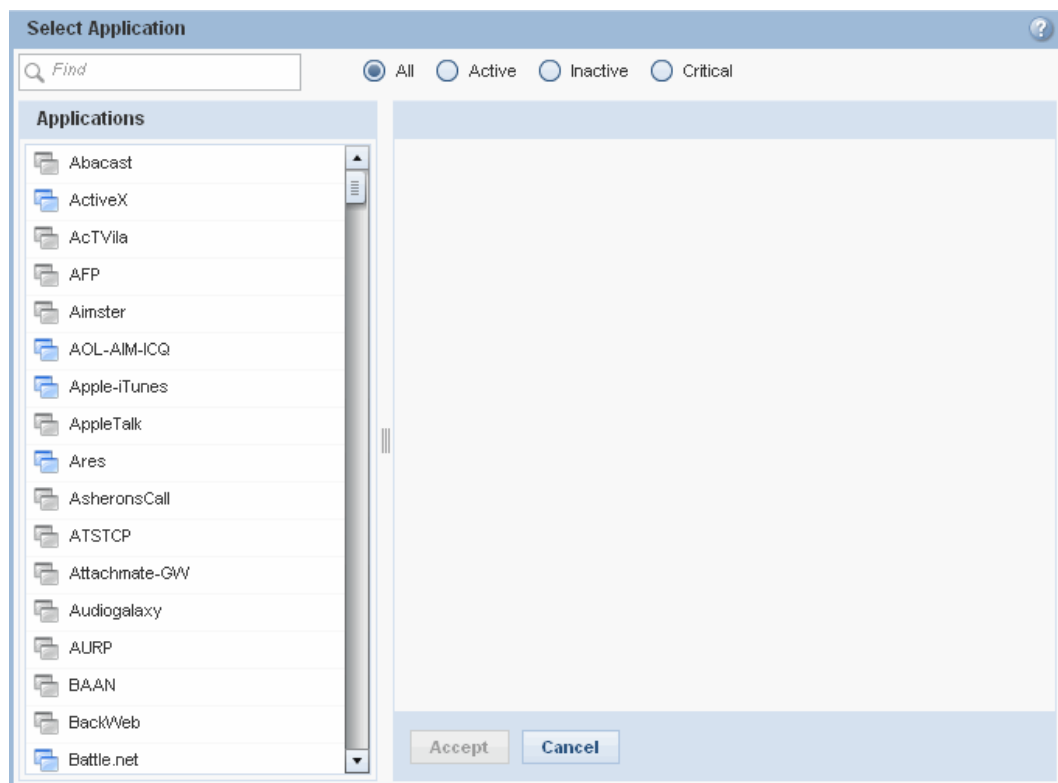
1. Click **Monitor**.
2. Select the portlet view that contains the Application Performance portlet that you want to configure.
3. Click the  icon in the portlet's title bar and select **Detailed Options** from the pop-up menu. The **Details** tab is displayed.
4. Click **Add**. A new application definition form displays at the top of the *Performance Configurations* section of the screen.



The **Performance Configurations** dialog box is shown. It includes a dropdown menu for selecting an application, a **Select...** button, and a **Remove** button. Below these are three performance level thresholds: **Normal** (0 to 100, currently at 100%), **Warning** (0 to 100, currently at 75%), and **Critical** (0 to 100, currently at 50%). To the right, a table lists parameters with their baseline values and weights.

Parameter	Baseline	Weight(%)
Loss:	0.001	30%
Availability:	100 %	40%
Latency:	30 ms	30%

5. In the **Application** field click **Select**. The *Application Selector* dialog box is displayed.



The **Select Application** dialog box is shown. It features a search bar labeled *Find* and four radio buttons for filtering: **All** (selected), **Active**, **Inactive**, and **Critical**. On the left is a list of applications, and on the right is a large empty area for details. At the bottom are **Accept** and **Cancel** buttons.

Applications

- Abacast
- ActiveX
- AcTVila
- AFP
- Aimster
- AOL-AIM-ICQ
- Apple-iTunes
- AppleTalk
- Ares
- Asheron'sCall
- ATSTCP
- Attachmate-GW
- Audiogalaxy
- AURP
- BAAN
- BackWeb
- Battle.net

6. Select the application you want to monitor in the *Applications* pane and click **Accept**. Note that the application you want to add must already be defined in IntelligenceCenter. If the application is not listed, you must [create an application](#) definition for it.

Note: Although both TCP and UDP applications are listed, only TCP applications will provide meaningful results in the Application Performance portlet.

7. Specify the **Performance Level Threshold** percentages for the **Critical**, **Warning**, and **Normal** levels by sliding the lever to the left to decrease the percentage or to the right to increase the percentage. The thresholds are used to set the boundaries for an application performance violation.

If the application's health percentage is between **Warning** and **Normal** levels, the application will be considered to be in compliance of the application performance thresholds you defined (Normal level, green). If the percentage is between **Critical** and

Warning levels, the application will be considered to be at a Warning level (orange). If the percentage is below **Critical** level, the application will be considered to be in violation of the application performance threshold (Critical level, red).


8. Specify the baseline weight and percentage values for the **Availability**, **Latency**, and **Loss** parameters.
 - To specify weights, enter a value in the text box next to each parameter. The weights you assign to each parameter must add up to 100. To determine the weight, consider how much each of these factors contributes to the application's health.
 - To set the baseline percentages for the parameters, slide the lever to the left to decrease the value or to the right to increase the value. The baseline value represents the desired or expected value for the parameter. As part of the calculation of an application's health, each parameter's current value is compared to its baseline value. Use the following table as a guideline in tuning the parameter values for an application.

Parameter	Tuning Considerations
Loss	By default, the baseline value for loss is 0.001 (.1%). To determine what the baseline should be for your application, start by sending a large number of pings and noting the drop rate. You can also gather more information by measuring the ME data on the PacketShaper appliances in your core and edge networks. The weight for this parameter is highly application dependent. Some applications, such as VoIP, are not adversely affected by reasonable amounts of loss, while other applications are extremely sensitive to loss. The default weight for loss is 30%.
Latency	By default, the baseline value for latency is 30ms. To determine what the baseline should be for your application, start by examining the ping times from client systems to server systems on your network, particularly with large payload pings. The weight for this parameter is highly application dependent. Many applications perform well with reasonably high latencies as long as there is little loss. The default weight for latency is 30%.
Availability	The default value of 100% is a good baseline for this parameter because you should expect your servers to be online and accessible at all times. If you have scheduled maintenance windows or other factors which predictably reduce the time the server will be available, you might want to lower this baseline. You would typically want to set the weight of this parameter fairly high because it measures the uptime of the servers and other network resources. The default weight of Availability is 40%.

9. To add additional applications to monitor, repeat steps 5-9.
10. When finished, click **Apply Changes**.

Modify an Existing Application Definition on the Application Performance Portlet

To modify an existing application definition on the Application Performance portlet:

1. Click **Monitor**.
2. Select the portlet view that contains the Application Performance portlet that you want to configure.
3. Click the  icon in the portlet's title bar and select **Detailed Options** from the pop-up menu. The **Details** tab is displayed.
4. Locate the application definition that you want to modify.
5. Specify the **Performance Level Threshold** percentages for the **Critical**, **Warning**, and **Normal** levels by sliding the lever to the left to decrease the percentage or to the right to increase the percentage. The thresholds are used to set the boundaries for an application performance violation.

If the application's health percentage is between **Warning** and **Normal** levels, the application will be considered to be in compliance of the application performance thresholds (Normal level, green). If the percentage is between **Critical** and **Warning** levels, the application will be considered to be at a Warning level (Warning level, orange). If the percentage is below **Critical** level, the application will be considered to be in violation of the application performance thresholds (Critical level, red).


6. Specify the baseline weight and percentage values for the **Availability**, **Latency**, and **Loss** parameters.
 - To specify weights, enter a value in the text box next to each parameter. The weights you assign to each parameter must add up to 100. To determine the weight, consider how much each of these factors contributes to the application's health.
 - To set the baseline percentages for the parameters, slide the lever to the left to decrease the value or to the right to increase the value. The baseline value represents the desired or expected value for the parameter. As part of the calculation of an application's health, each parameter's current value is compared to its baseline value. Use the following table as a guideline in tuning the parameter values for an application.

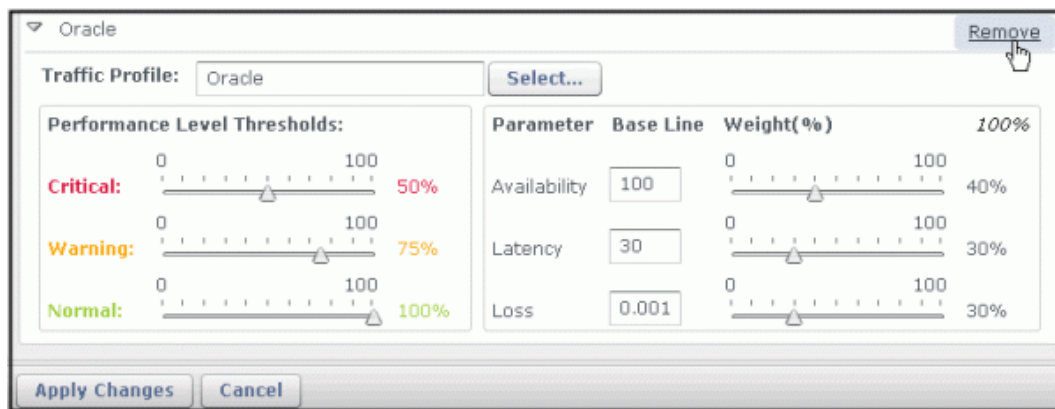
Parameter	Tuning Considerations
Loss	By default, the baseline value for loss is 0.001 (.1%). To determine what the baseline should be for your application, start by sending a large number of pings and noting the drop rate. You can also gather more information by measuring the ME data on the PacketShaper appliances in your core and edge networks. The weight for this parameter is highly application dependent. Some applications, such as VoIP, are not adversely affected by reasonable amounts of loss, while other applications are extremely sensitive to loss. The default weight for loss is 30%.
Latency	By default, the baseline value for latency is 30ms. To determine what the baseline should be for your application, start by examining the ping times from client systems to server systems on your network, particularly with large payload pings. The weight for this parameter is highly application dependent. Many applications perform well with reasonably high latencies as long as there is little loss. The default weight for latency is 30%.
Availability	The default value of 100% is a good baseline for this parameter because you should expect your servers to be online and accessible at all times. If you have scheduled maintenance windows or other factors which predictably reduce the time the server will be available, you might want to lower this baseline. You would typically want to set the weight of this parameter fairly high because it measures the uptime of the servers and other network resources. The default weight of Availability is 40%.

7. To customize other applications, repeat steps 4-6.
8. When finished, click **Apply Changes**.

Remove an Application Definition from the Application Performance Portlet

To remove an application definition from the Application Performance portlet:

1. Click **Monitor**.
2. Select the portlet view that contains the Application Performance portlet that you want to configure.
3. Click the  icon in the portlet's title bar and select **Detailed Options** from the pop-up menu. The **Details** tab is displayed.
4. Locate the application definition that you want to delete.
5. Click the **Remove** link. (You may need to scroll to the right or widen the pane to see the link).



6. When prompted to confirm the delete operation, click **Yes**.
7. Click **Apply Changes**.

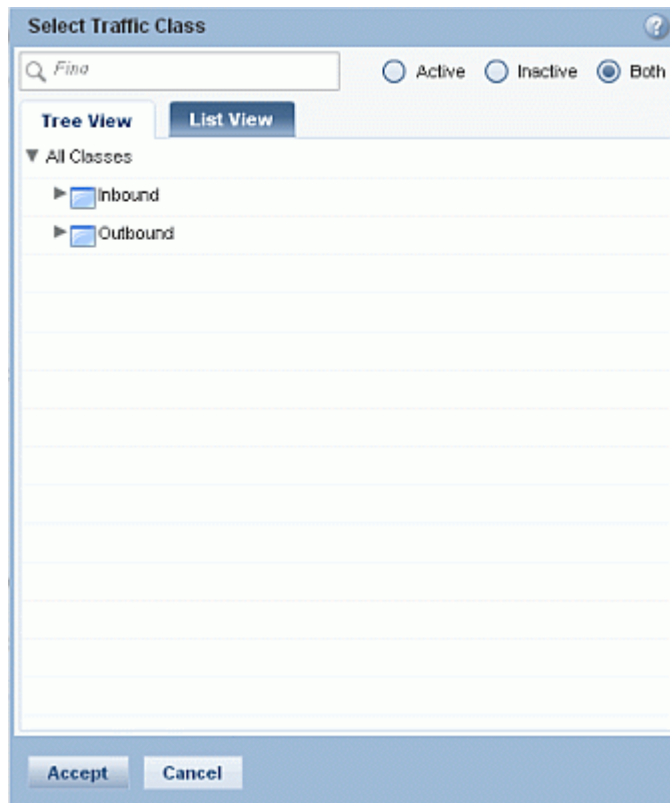
Configure the Class Utilization Portlet

The Class Utilization portlet provides at-a-glance monitoring of the traffic on your network. With each instance of the Class Utilization portlet that you configure, you can monitor and compare the average and peak bandwidth usage of up to five traffic classes. This allows you to group classes for specific monitoring purposes.

Before you can begin monitoring traffic classes, you must configure the Class Utilization portlet to show the classes you want to compare/monitor. Keep in mind that you can create several different instances of the Class Utilization portlet, each showing a different set of classes.

1. From within IC, click the **Monitor** tab.
2. Drag the **Class Utilization** portlet from the Portlets pane into the selected portlet view.
3. To define the traffic classes to monitor, click **Add Class**. The Add Traffic Class dialog box displays.

4. Specify the network group, subgroup, device or view for which the portlet should display data:
 - a. In the **Network Group or Device** field, click **Select**. The Select Group or View dialog box displays.
 - b. Select the network group, subgroup, or device from the **Devices** tab or select the logical network view on the **Views** tab for which you want to monitor class data and then click **Accept**.
5. Add the traffic classes you want to compare/monitor:
 - a. In the **Traffic Class** field, click **Select**. The Select Traffic Class dialog box displays.





- b. Select a traffic class:
 - To display an alphabetical listing of traffic classes, including path, select the **List View** tab.
 - To display the traffic class tree, select **Tree View**. In this view you can expand the parent class for the class you want to select; just click the arrow icons next to each class. The down arrow ▼ icon indicates that the class is expanded; the right arrow ► icon indicates that the class is collapsed.
 - To narrow the list of displayed classes, select the **Active** or **Inactive** radio button. By default, **Both** active and inactive classes are displayed.
 - After you select the traffic class, click **Accept**. The class you selected appears in the **Class list for comparison** list.
 - c. If you want to add another traffic class, repeat Steps a and b. You can add up to five traffic classes.
 - d. When you have added all of the traffic classes that you want to compare, click **Save**.
6. Select the time range for which you want to display data by clicking one of the following time span buttons at the top of the portlet:



Time Range	Description
1H	<p>Shows data for the last sixty minutes up to the last quarter hour. For example, if the portlet refreshes at 11:20, the portlet will display data collected between 10:20 and 11:15. This is the default time range.</p> <p>The units that are displayed depend on the granularity settings for the DataCollector. If the granularity is set to one minute (the default); the portlet will plot a data point for every minute. Similarly, if you set the granularity to 5 minutes, the portlet will plot the data points at 5-minute intervals.</p>
TDY	<p>Shows all data collected today, from midnight up through last full quarter hour. For example, if the portlet refreshes at 11:20, the portlet will display data collected between 12:00 AM (midnight) and 11:15 AM.</p> <p>The units that are displayed depend on the granularity settings for the DataCollector. If the granularity is set to one minute (the default); the portlet will plot a data point for every minute. Similarly, if you set the granularity to 5 minutes, the portlet will plot the data points at 5-minute intervals.</p>
YD	<p>Shows all data collected yesterday, beginning at midnight yesterday and ending at midnight today.</p> <p>The units that are displayed depend on the granularity settings for the DataCollector. If the granularity is set to one minute (the default); the portlet will plot a data point for every minute. Similarly, if you set the granularity to 5 minutes, the portlet will plot the data points at 5-minute intervals.</p>
1W	<p>Shows data in one-hour units for the last seven day period up through the last full hour. For example, if you refresh the portlet at 11:20 on Tuesday, March 15, the portlet will display data collected between 11:00 on Tuesday, March 8 through 11:00 on Tuesday, March 15.</p>
1M	<p>Shows data in one-hour units for the last month up through the last full hour. For example, if you refresh the portlet at 10:50 on March 15, the portlet will display data collected between 10:00 on February 15 through 10:00 on March 15.</p>

3M	Shows data in one-day units for the last three months up through the last full day. For example, if you refresh the portlet at 10:50 on Tuesday, March 15, 2011 the portlet will display data collected between 12:00 (midnight) on December 15, 2010 through 12:00 (midnight) on March 15, 2011.
6M	Shows data in one-day units for the last six months up through the last full day. For example, if you refresh the portlet at 10:50 on Tuesday, March 15, 2011 the portlet will display data collected between 12:00 (midnight) on September 15, 2010 through 12:00 (midnight) on March 15, 2011.
1Y	Shows data in one-day units for the last year up through the last full day. For example, if you refresh the portlet at 10:50 on Tuesday, March 15, 2011 the portlet will display data collected between 12:00 (midnight) on March 15, 2010 through 12:00 (midnight) on March 15, 2011.

7. (optional) Modify which graphs are displayed. By default, the portlet shows graphs for both average rate and peak rate. However, if you are only interested in one of the graphs, you can remove the other one as follows:
 - a. Click the  icon in the portlet's title bar and select **Details** from the pop-up menu. The **Details** configuration tab appears.
 - b. In the **Select Line Chart** field, deselect the graph you do not want to display on the portlet.
 - c. Click **Apply Changes**.
8. (optional) To configure the portlet to refresh automatically:
 - a. Click the  icon in the portlet's title bar and select **Configuration** from the pop-up menu. The **General** configuration tab appears.
 - b. Check the **Automatically update every xx minutes** checkbox. When you enable automatic update, the portlet will automatically update at a fixed interval depending on which time range you select when you view the portlet:
 - If you select **Last 60 minutes, Today, or Yesterday (1H, TDY, or YD)** if you set the time range from the [data view](#) as the time range, the portlet will automatically refresh every 15 minutes.
 - If you select **Last seven days (1W)** in the [data view](#) as the time range, the portlet will automatically refresh every hour.
 - If you select a time range of 1M or greater, the portlet will automatically refresh every 1440 minutes (24 hours).
 - c. Click **Apply Changes**.

Configure the Network Efficiency Portlet

Before you can begin monitoring network efficiency, you must configure the Network Efficiency portlet to show the classes you want to compare/monitor. Keep in mind that you can create several different instances of the Network Efficiency portlet, each showing a different set of classes.

1. From within IC, click the **Monitor** tab.
2. Drag the Network Efficiency portlet from the Portlets pane into the selected portlet view.
3. To define the traffic classes to monitor, click **Add Class**. The Add Traffic Class dialog box displays.

Add Traffic Class

Network Group or Device **Select...**


Traffic Class **Select...**

**You can select maximum 5 classes.*

Class list for comparison

Save **Cancel**

4. Specify the network group, subgroup, device or view for which the portlet should display data:
 - a. In the **Network Group or Device** field, click **Select**. The Select Group or View dialog box displays.
 - b. Select the network group, subgroup, or device from the **Devices** tab or select the logical network view on the **Views** tab for which you want to monitor class data and then click **Accept**.
5. Add the traffic classes you want to compare/monitor:
 - a. In the **Traffic Class** field, click **Select**. The Select Traffic Class dialog box displays.



b. Select a traffic class:


- To display an alphabetical listing of traffic classes, including path, select the **List View** tab.
- To display the traffic class tree, select **Tree View**. In this view you can expand the parent class for the class you want to select; just click the arrow icons next to each class. The down arrow ▼ icon indicates that the class is expanded; the right arrow ► icon indicates that the class is collapsed.
- To narrow the list of displayed classes, select the **Active** or **Inactive** radio button. By default, Both active and inactive classes are displayed.
- After you select the traffic class, click **Accept**. The class you selected appears in the **Class list for comparison** list.

c. If you want to add another traffic class, repeat Steps a and b. You can add up to five traffic classes.

Note: If you select a non-TCP traffic class, the corresponding graph will display a network efficiency of 100%.


d. When you have added all of the traffic classes that you want to compare, click **Save**.

6. Select the time range for which you want to display data by clicking one of the following time span buttons at the top of the portlet:



Time Range	Description
1H	<p>Shows data for the last sixty minutes up to the last quarter hour. For example, if the portlet refreshes at 11:20, the portlet will display data collected between 10:20 and 11:15. This is the default time range.</p> <p>The units that are displayed depend on the granularity settings for the DataCollector. If the granularity is set to one minute (the default); the portlet will plot a data point for every minute. Similarly, if you set the granularity to 5 minutes, the portlet will plot the data points at 5-minute intervals.</p>
TDY	<p>Shows all data collected today, from midnight up through last full quarter hour. For example, if the portlet refreshes at 11:20, the portlet will display data collected between 12:00 AM (midnight) and 11:15 AM.</p> <p>The units that are displayed depend on the granularity settings for the DataCollector. If the granularity is set to one minute (the default); the portlet will plot a data point for every minute. Similarly, if you set the granularity to 5 minutes, the portlet will plot the data points at 5-minute intervals.</p>
YD	<p>Shows all data collected yesterday, beginning at midnight yesterday and ending at midnight today.</p> <p>The units that are displayed depend on the granularity settings for the DataCollector. If the granularity is set to one minute (the default); the portlet will plot a data point for every minute. Similarly, if you set the granularity to 5 minutes, the portlet will plot the data points at 5-minute intervals.</p>
1W	<p>Shows data in one-hour units for the last seven day period up through the last full hour. For example, if you refresh the portlet at 11:20 on Tuesday, March 15, the portlet will display data collected between 11:00 on Tuesday, March 8 through 11:00 on Tuesday, March 15.</p>
1M	<p>Shows data in one-hour units for the last month up through the last full hour. For example, if you refresh the portlet at 10:50 on March 15, the portlet will display data collected between 10:00 on February 15 through 10:00 on March 15.</p>

3M	Shows data in one-day units for the last three months up through the last full day. For example, if you refresh the portlet at 10:50 on Tuesday, March 15, 2011 the portlet will display data collected between 12:00 (midnight) on December 15, 2010 through 12:00 (midnight) on March 15, 2011.
6M	Shows data in one-day units for the last six months up through the last full day. For example, if you refresh the portlet at 10:50 on Tuesday, March 15, 2011 the portlet will display data collected between 12:00 (midnight) on September 15, 2010 through 12:00 (midnight) on March 15, 2011.
1Y	Shows data in one-day units for the last year up through the last full day. For example, if you refresh the portlet at 10:50 on Tuesday, March 15, 2011 the portlet will display data collected between 12:00 (midnight) on March 15, 2010 through 12:00 (midnight) on March 15, 2011.



5. (optional) To configure the portlet to refresh automatically:
 - a. Click the  icon in the portlet's title bar and select **Configuration** from the pop-up menu. The **General** configuration tab appears.
 - b. Check the **Automatically update every xx minutes** checkbox. When you enable automatic update, the portlet will automatically update at a fixed interval depending on which time range you select when you view the portlet:
 - If you select **Last 60 minutes**, **Today**, or **Yesterday** (**1H**, **TDY**, or **YD** if you set the time range from the [data view](#)) as the time range, the portlet will automatically refresh every 15 minutes.
 - If you select **Last seven days** (**1W** in the [data view](#)) as the time range, the portlet will automatically refresh every hour.
 - If you select a time range of **One month** (**1M** in the [data view](#)) or greater, the portlet will automatically refresh every 1440 minutes (24 hours).
8. Click **Apply Changes**. The [data view](#) displays.

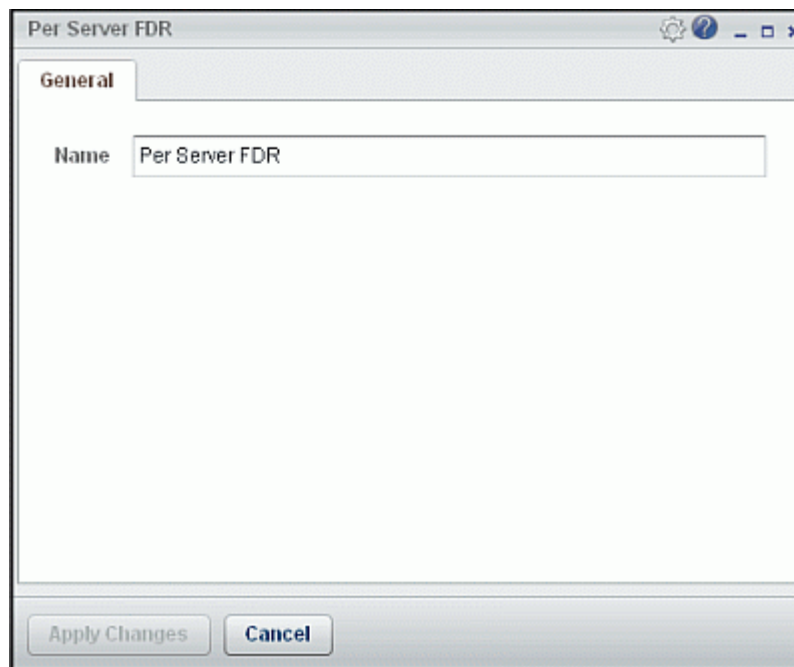
Configure the Per Server FDR Portlet

Monitor Servers With the Per Server FDR Portlet


The [Per Server FDR portlet](#) allows you to perform intruder detection monitoring on the servers within your network and allows you to flag any services that you think are suspicious.

To configure the Per Server FDR portlet:

1. Click **Monitor**.
2. Select the portlet view that contains the Per Server FDR portlet that you want to configure or [drag a new Per Server FDR portlet](#) from the **Portlets** pane to the portlet view.
3. To configure the portlet, click the  icon in the portlet's title bar and select **Details** from the pop-up menu.
4. Specify the network group for which the portlet should display data by clicking **Select** in the **Group** field. The *Select Group* dialog box is displayed. Select the network group or sub-group and then click **Accept**. Click **Apply Changes** to save your selection. The portlet will not display any data until you select a **Group**.
5. If you want to specify a new name to display for the portlet instance, select the **General** tab, enter a **Name** and then click **Apply Changes**.
6. To return to the data view, click the  icon in the portlet's title bar and select **View Data** from the pop-up menu.



7. To monitor a specific server, select the server IP address from the **Server** drop-down list. Note that if you have multiple servers that you want to routinely monitor, you can open separate portlet instances for each one.
8. Select the time range over which to monitor. By default, the portlet tracks all inbound flows to the server in the last 15 minutes (rounded to the last quarter hour). However, you can select a different time range from the drop-down menu. After you select a time range, an updated list of services that have had flows to the server during the selected time

range is displayed on the **Summary** tab. For each service, the number of flows, bytes, and packets is displayed. Initially each service is marked with a warning  icon.














Per Server FDR

Server: **10.9.50.93** Last 60 minutes Refresh

Updated Nov 17, 2008 4:00 PM

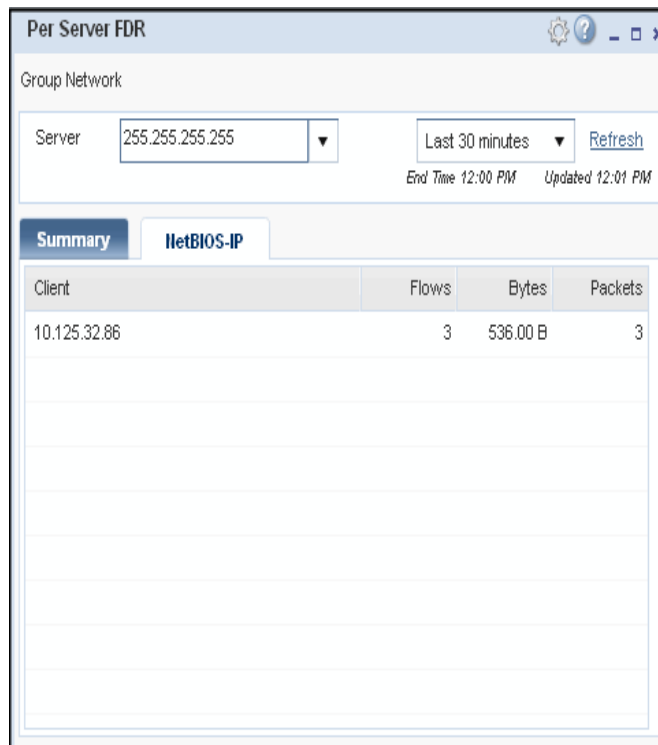
Summary

☒ Abnormal ☐ All

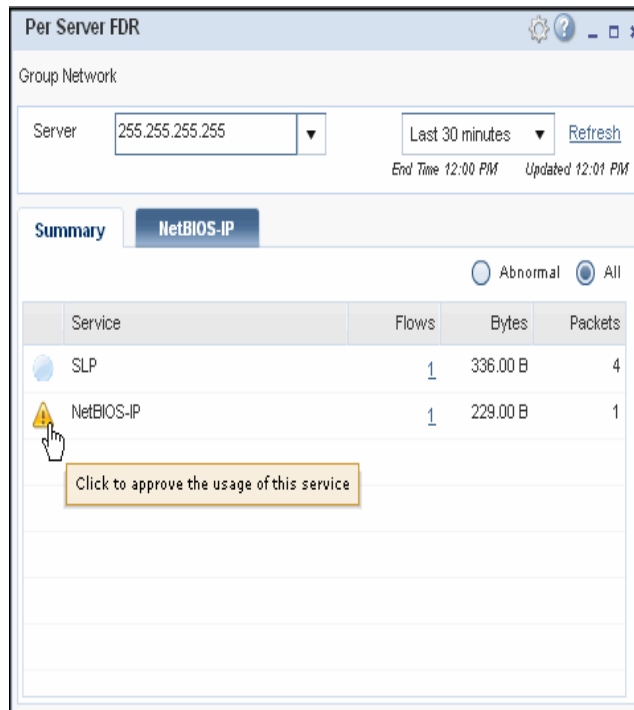
	Service	Flows	Bytes	Packets
	Palo_Alto/HTTP	120	4.25 MB	4,067
	Palo_Alto/DCOM	106	29.15 KB	551
	Palo_Alto/DNS	87	36.32 KB	123
	Palo_Alto/MS-Exchange	70	16.62 KB	319
	Palo_Alto/Default	23	25.04 KB	112
	Palo_Alto/SSL	11	1.41 MB	1,195
	Palo_Alto/Skype	10	85.61 KB	560
	Palo_Alto/GoogleVideo	9	4.58 KB	54
	Palo_Alto/CMP	6	576.00 B	10
	Palo_Alto/SOAP-HTTP	4	3.90 KB	20
	Palo_Alto/YouTube	3	8.86 MB	6,222
	Palo_Alto/FlashVideo	2	14.34 MB	10,053
	Palo_Alto/CIFS	2	7.02 KB	22

Click to approve the usage of this service

- To display more information about a particular service, click the number in the corresponding **Flows** column of the portlet to drill down to the actual flow detail records (FDR). A second tab opens in the portlet window with additional information about the specific hosts that originated the traffic for the selected service.



10. For any service that you determine to be normal, go back to the **Summary** tab and click the warning ⚠ icon to approve the service. The icon changes to the approved traffic 🟢 icon.



11. After you designate services as approved or disapproved you can filter the list of services to show either **Abnormal** traffic or **All** traffic.
12. To update the data that is displayed on the portlet, click **Refresh**.

Predefined Per Server FDR Portlet Time Ranges

Because the data that is displayed in the portlet is based on data that is collected and stored in time-based database tables, the built-in time ranges are based on standard hour, half hour, and quarter hour intervals (for example, 9:00, 9:15, 9:30) and are rounded to the nearest quarter hour. Therefore, the actual time range that is used depends on the time at which you refresh the portlet data. For example, if you refresh the portlet at 11:20, the end of the range will be set to 11:15, the last full quarter hour.

Note: The Per Subnet FDR portlet and Application Performance portlet use different predefined time ranges.


You can select one of the following time ranges when running the Per Server FDR portlet:

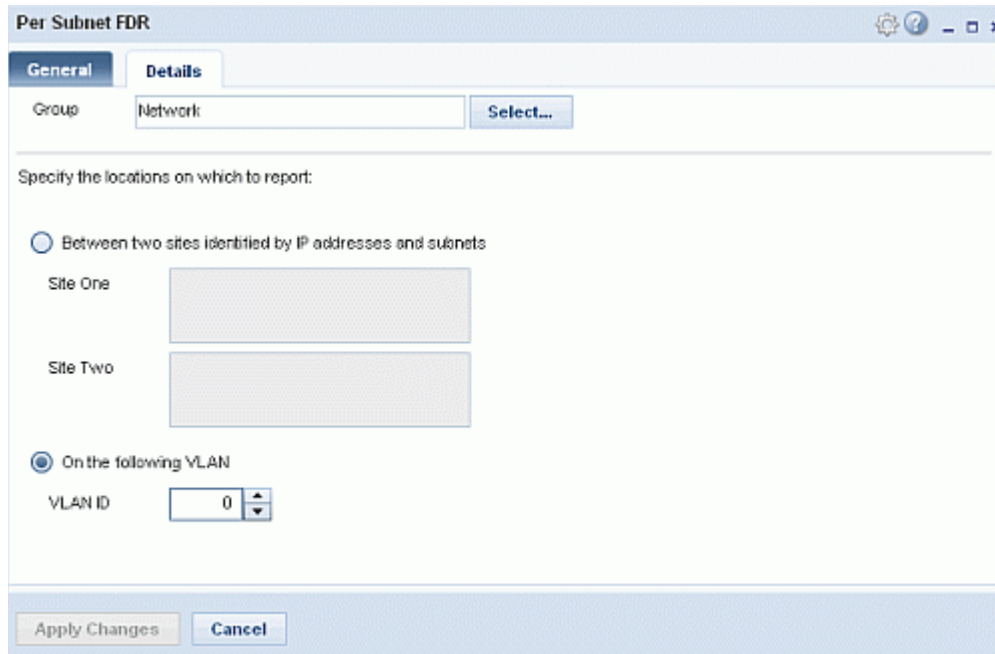
Time Range	Description
Last 15 minutes	Shows data for the last 15 minutes, ending at the last quarter hour. For example, if you refresh the portlet at 11:20, this selection would show data collected between 11:00 and 11:15. This is the default time range for this portlet.
Last 30 minutes	Shows data for the last full 30-minute time interval. For example, if you refresh the portlet at 11:20, the portlet will display data collected between 10:45 and 11:15.
Last 60 minutes	Shows data for the last full 60-minute interval. For example, if you refresh the portlet at 11:20, the portlet will display data collected between 10:15 and 11:15.
Last 4 hours	Shows data for the last full 4 hours. For example, if you refresh the portlet at 11:20, the portlet will display data from 7:15 to 11:15.
Last 8 hours	Shows data for the last full 8 hours. For example, if you refresh the portlet at 11:20, the portlet will display data from 3:15 to 11:15.

Configure the Per Subnet FDR Portlet

The [Per Subnet FDR portlet](#) displays statistics for traffic between two sites (such as two subnets) or on a specific virtual LAN (VLAN).

To configure the Per Subnet FDR portlet:

1. Click **Monitor**.
2. Select the portlet view that contains the Per Subnet FDR portlet that you want to customize or [drag a new Per Subnet FDR portlet](#) from the **Portlets** pane to the portlet view.
3. To configure the portlet, click the  icon in the portlet's title bar and select **Details** from the pop-up menu.



4. Specify the network group or view for which the portlet should display data by clicking **Select** in the **Group** field. The *Select Group or View* dialog box is displayed. Select the network group, sub-group or device on the **Devices** tab or the logical network view on the **Views** tab to which to restrict this portlet and then click **Accept**. The portlet will not display any data until you select a **Group**.
5. If you want to configure the sites or VLANs to monitor, choose one of the following:
 - o Choose **Between two sites identified by IP addresses and subnets** and then define the address space associated with **Site One** and **Site Two**. Each item should be separated by a comma. Enter one or more of the following: IP address (x.x.x.x where 0<x<255), IP address range (x.x.x.x-x.x.x.y), subnet (x.x.x.x/y). This field allows you to create a discontinuous address space on the network.
-or-
 - o Choose **On the following VLAN** and then select the ID number of any 802.1q VLAN (the range is 0-4095).
6. Click **Apply Changes** to save your settings.
7. If you want to specify a new name to display for the portlet instance or modify the time range over which to monitor, select the **General** tab.

Per Subnet FDR

General Details

Name Per Subnet FDR

Select Time Range:

Last hour

☒ Automatically update every 15 minutes

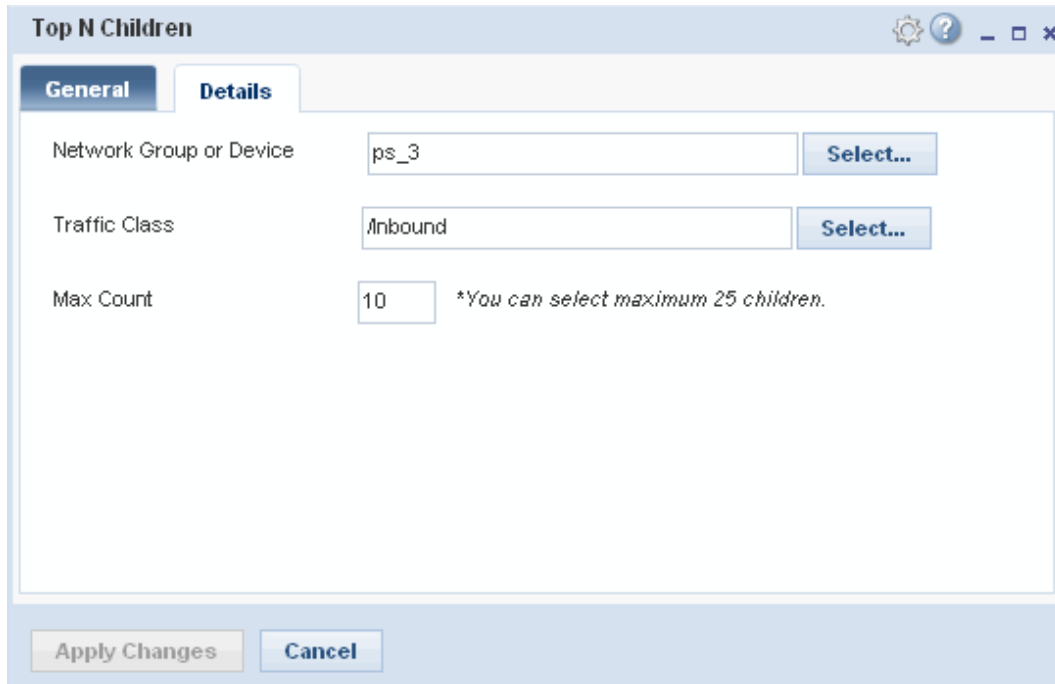
Apply Changes Cancel

8. Enter a **Name** to display for the portlet instance.
9. Select the time range over which to monitor. By default, the portlet monitors application traffic collected today (rounded to the last full hour). However, you can select a different predefined time range from the drop-down menu or select **Other** to define your own time range. For best performance, you should use a predefined time range whenever possible.
10. If you want the portlet to refresh automatically, check the **Automatically update every xx minutes** checkbox and then set the frequency at which you want automatic updates to occur. The range is 15 minutes (the default) up to 1200 minutes.
11. Click **Apply Changes**.

Configure the Top N Children Portlet

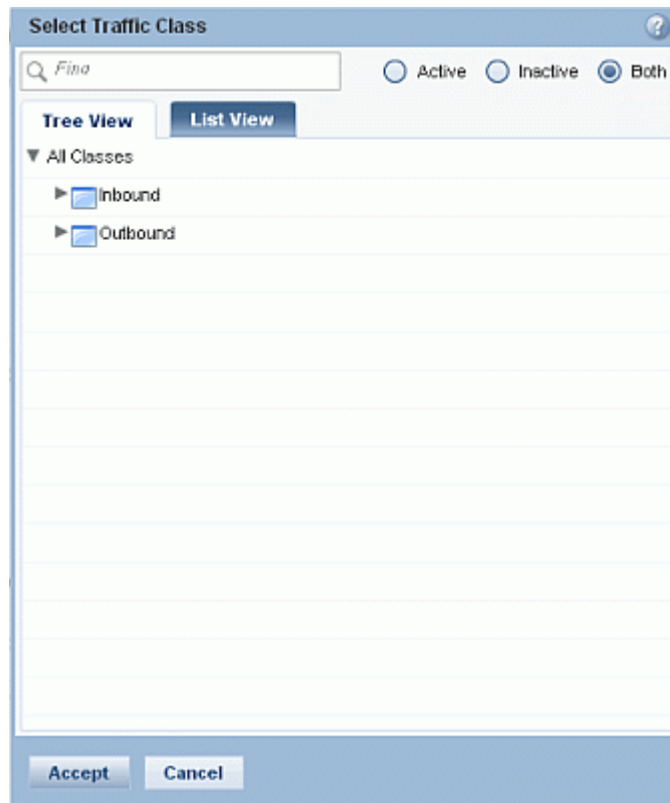
To configure the Top N Children portlet:

1. Click **Monitor**.
2. Select the portlet view that contains the Top N Children portlet that you want to customize or [drag a new Top N Children portlet](#) from the **Portlets** pane to the portlet view.
3. To configure the portlet, click the  icon in the portlet's title bar and select **Details** from the pop-up menu.



The screenshot shows the 'Top N Children' configuration dialog box. It has a title bar with a gear icon, a question mark icon, and standard window controls. Below the title bar are two tabs: 'General' and 'Details'. The 'Details' tab is selected. The dialog contains three input fields with corresponding 'Select...' buttons: 'Network Group or Device' with the value 'ps_3', 'Traffic Class' with the value 'Inbound', and 'Max Count' with the value '10'. A note next to the 'Max Count' field states '*You can select maximum 25 children.' At the bottom of the dialog are two buttons: 'Apply Changes' and 'Cancel'.

4. Specify the network group or view for which the portlet should display data by clicking **Select** in the **Group** field. The *Select Group or View* dialog box is displayed. Select the network group, sub-group or device on the **Devices** tab or the logical network view on the **Views** tab to which to restrict this portlet and then click **Accept**. The portlet will not display any data until you select a **Group**.
5. Add the traffic class for which you want to monitor child classes by clicking **Select** in the **Traffic Class** field. The *Select Traffic Class* dialog box displays.




- a. Select a traffic class:
 - To display an alphabetical listing of traffic classes, including path, select the **List View** tab.
 - To display the traffic class tree, select **Tree View**. In this view you can expand the parent class for the class you want to select; just click the arrow icons next to each class. The down arrow ▼ icon indicates that the class is expanded; the right arrow ► icon indicates that the class is collapsed.
 - To narrow the list of displayed classes, select the **Active** or **Inactive** radio button. By default, **Both** active and inactive classes are displayed.
 - b. After you select the traffic class, click **Accept**. The class you selected appears in the **Class list for comparison** list.
6. Specify the number of child classes to monitor by entering a value in the **Max Count** field (up to a maximum of 25). The portlet will then display the top bandwidth consuming traffic classes up to the limit you specify. It will then aggregate data for the remaining classes into a group called *Others*.
 7. Select the **General** tab.

8. If you want to customize the portlet name, enter a new **Name**. This is a good idea if you plan to monitor multiple classes in separate portlet instances, you may want to name each portlet instance according to the class you are monitoring. For example, you might name one instance *Top 10 Children for /Inbound* and one instance *Top 10 Children for /Outbound*.
9. Select the time range over which to monitor. By default, the portlet monitors data collected over the last 60 minutes. However, you can select a different pre-defined time range from the **Select Time Range** drop-down.
10. (optional) To configure the portlet to refresh automatically check the **Automatically update every xx minutes** checkbox. When you enable automatic update, the portlet will automatically update at a fixed interval depending on which time range you select when you view the portlet:
 - If you select **Last 60 minutes**, **Today**, or **Yesterday** as the time range, the portlet will automatically refresh every 15 minutes.
 - If you select **Last seven days** as the time range, the portlet will automatically refresh every hour.
 - If you select a time range of **One month** or greater, the portlet will automatically refresh every 1440 minutes (24 hours).
11. Click **Apply Changes**. The [data view](#) displays.

Configure the VoIP Performance Portlet

Before you can use the VoIP Performance portlet to monitor your VoIP traffic, you must configure it as follows:

1. From within IntelligenceCenter, click the **Monitor** tab.
2. Drag the **VoIP Performance** portlet from the Portlets pane into the selected portlet view. If the VoIP Performance portlet is not displayed on the Portlets pane, you may not have successfully plugged it in to IC yet.
3. Click the  icon in the portlet's title bar and select **Configuration** from the pop-up menu. The **General** configuration tab appears.
4. If you want to customize the portlet name, enter a new **Name**.
5. Select the time range over which to monitor VoIP traffic. By default, the portlet monitors traffic collected today (from midnight up through the last full hour). However, you can select from the following time ranges:


Time Range	Description
Today	Shows data from midnight today up through last full hour. For example, if the portlet refreshes at 11:20, the portlet will display data collected between 12:00 AM (midnight) and 11:00 AM. This is the default time range.
Last 24 hours	Reports data for the last twenty-four hour period. For example, if the report runs at 11:20 on February 3, 2011, this selection would show data collected between 11:00:00 on February 2, 2011 and 11:00:00 on February 3, 2011.

6. If you want the portlet to refresh automatically every 15 minutes, check the **Automatically update every xx minutes** checkbox.
7. To configure the specific network group and application to monitor, select the **Details** tab.
8. Specify the network group or view for which the portlet should display data by clicking **Select** in the **Group** field. The *Select Group or View* dialog box is displayed. Select the network group, sub-group or device on the **Devices** tab or the logical network view on the **Views** tab to which to restrict this Application Performance portlet and then click **Accept**.
9. To change the number of flow records that are displayed on the portlet, select a value from the **Maximum call records returned** field. The range is 1-5000; 200 is the default.
10. Select the **Codec** that the VoIP application you are monitoring uses. Each VoIP application, device, or service uses a specific codec (coder/decoder) for the conversion between analog and digital signals. To determine which codec to use, refer to the documentation provided by the VoIP application, device, or service provider. Selecting the correct codec ensures the most accurate reporting and calculation of VoIP metrics such as percent loss, latency, jitter, and MOS.
11. Specify whether to include or exclude latency from the MOS calculation. Latency values can only be calculated if you have a PacketShaper at each end of the VoIP call; if the VoIP traffic goes through a single PacketShaper only, the latency value will be zero. However, because MOS is calculated using latency values, latency values of zero may throw off the MOS calculation. You can configure the portlet so that latency values of zero are:

- **Excluded** from the MOS calculation. This is the best option to use if most of your VoIP traffic has non-zero values (that is, you have PacketShapers at each end of your VoIP calls).
 - **Included with a value of x ms.** By default, latency values of 0 are included in the MOS calculation. However, you can change the value of this field so that zero values are substituted with a different value (in the range of 0 to 9999). You may want to supply a non-zero value for this field if most of your VoIP traffic does not have latency values (that is, you do not have PacketShapers at each end of your VoIP calls) and you have consistent latency values on your network. In this case, use a value that represents a typical latency value on your network. If you substitute your own latency values, they will display in blue when displayed in flow details on the portlet.
12. When you are finished configuring the portlet, click **Apply Changes**. The [data view](#) displays.

View a Portlet


To view a portlet:

1. Click **Monitor**.
2. Select the portlet view that contains the portlet that you want to view. All the [portletlets that you have added](#) the selected view are displayed.
3. If you want to ensure that you are viewing the most up-to-date data, click the  icon in the portlet's title bar and select **Refresh Data** from the pop-up menu.
4. The way you interact with a portlet, depends on the [portlet type](#).
5. You can also:
 - [Arrange how the portlets are displayed in the view](#)
 - [Customize the portlet's configuration settings](#)
 - [Delete the portlet](#)

NOTE: The portlets will not display any data until you [configure](#) them to monitor a specific network group, sub-group, device, or view for them to monitor.


Remove a Portlet

To remove a portlet:

1. Click **Monitor**.
2. Select the portlet view that contains the portlet you want to delete. For example, if you want to delete a portlet from the *My View* portlet view, click the **My View** tab.
3. Click the  in the top right-hand corner of the portlet window.
4. Click **Yes** when prompted to confirm the delete operation.

Arrange Portlet Windows

To arrange how portlet windows are displayed in the selected portlet view:

1. Click **Monitor**.
2. Select the portlet view that contains the portlet windows you want to arrange. For example, if you want to arrange the portlet windows on the *My View* portal pane, click the **My View** tab.
3. Click the **View options**  button and select one of the following options:
 - If you want to arrange the portlets from left to right and top to bottom, select **Tile Portlets**.
 - If you want to arrange the portlet windows from top to bottom, select **Tile Portlets Vertically**.
 - If you want to arrange the portlet windows from left to right, select **Tile Portlets Horizontally**.
 - If you want to stack the portlet windows so that the title bar of each shows, select **Cascade**.
4. If you want to manually arrange the portlet windows, select the title bar of each portlet and drag it to where you want it to go. If you want to turn on or off grid lines to help you arrange the portlets, select **Toggle Grids** from the **Options** drop-down menu.

Reports

IntelligenceCenter Reports

IntelligenceCenter provides a comprehensive set of reports that allow you monitor the performance of the applications, devices, and hosts that are running on your network. The data that IntelligenceCenter displays is collected from the PacketShaper appliances throughout your network and aggregated according to options you define when you [run a report](#). You can generate reports that include data for your entire network, or you can define a specific [geographic region or logical group](#) — such as a business unit — that is of interest to you. You can also generate reports for a specific [site](#) (subnet class).

The following sections describe the IntelligenceCenter report categories and list the reports that are available in each category. For detailed information about a specific report, click on the corresponding link.

Application Reports

Application reports give you insight into the performance of the [applications](#) on your network. Application reports are only available if your DataCollector is configured to collect [metric data](#).

- [Application Activity](#)
- [Application Response Time](#)
- [Top Applications](#)
- [Top Immediate Children](#)

Device Reports

Device reports detail the performance of specific PacketShaper appliances or groups of appliances. Device reports are only available if your DataCollector is configured to collect [metric data](#).

- [Device Compression](#)
- [Link Utilization](#)
- [TCP Health](#)
- [Top Traffic Classes Summary](#)
- [Traffic Class Compression](#)
- [Traffic Class Response Time](#)
- [Traffic Class Utilization](#)
- [VoIP Statistics](#)

Host Reports

Host reports help you identify the hosts that consume the most bandwidth — in terms of bytes or packets — on your network as a whole or in a specific geographic or logical group. Host reports are only available if your DataCollector is configured to collect [FDR data](#).

- [Host Pairs Activity](#)
- [Top DSCP](#)
- [Top Host Pairs](#)
- [Top Listeners](#)
- [Top Services](#)
- [Top Talkers](#)
- [Top VLAN](#)

Site Reports

When your PacketShapers are configured with site-based traffic trees, you can use IC's site reports to display data associated with each of these pre-defined sites. For example, a PacketShaper might classify the network traffic by subnet, with each location (subnet) having its own traffic class (Cupertino, Sunnyvale, San_Jose, and so forth). For more information about creating site-based traffic trees, refer to [Create a Location-Based Traffic Tree with Per-Location Applications](#) in PacketGuide for PacketWise.


Site reports require that you have [defined the sites](#) in IC and configured your DataCollector to collect [metric data](#) and [FDR data](#).


- [Site Response Time](#)
- [Top Applications by Site](#)
- [Top Host Pairs by Site](#)
- [Top Listeners by Site](#)
- [Top Services by Site](#)
- [Top Sites](#)
- [Top Talkers by Site](#)

Application Reports

Application Activity Report

The Application Activity report details the average and peak rates in bits per second (bps) for the traffic classes that comprise the selected [application](#) during the reporting period. This report displays separate tables and line graphs detailing inbound and outbound application utilization. For each traffic class (component), each table displays the total number of packets and bytes used by the traffic class, the efficiency level percentage, and the average and peak rates for the traffic classes that comprised the application traffic generated during the reporting period. You can re-sort the table data by clicking on **Component**, **Bytes**, or **Packets** in the table header. This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.





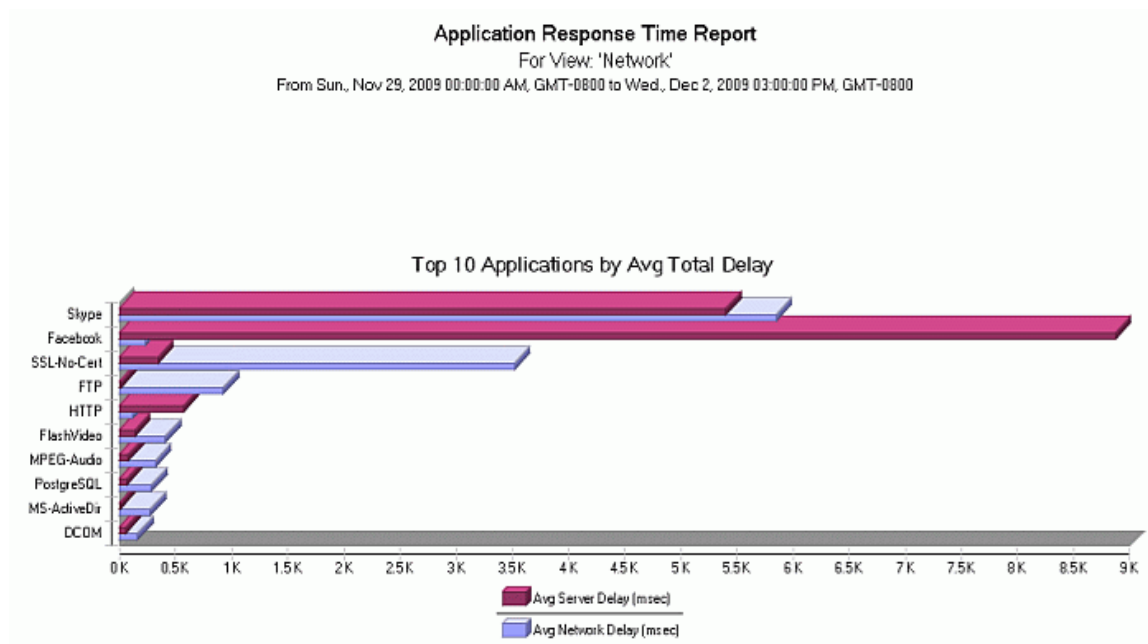
Application Response Time Report

Response time measurement (RTM) provides information about the amount of time connection-based TCP traffic spends traveling between a client and a server and the time used by the server itself. This allows you to investigate response times and identify the source of network delays. The Application RTM report contrasts RTM statistics for the applications that are running on your network. When you configure the report, you specify the network group, sub-group, device, or view to which to restrict the report. For example, you could restrict the report so that it shows application RTM statistics for a specific branch office or within a specific department.

This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.

Top-Level View

Application RTM Graph — Compares the average network delay and average server delay for the 10 applications with the highest average total delay in the selected network group during the reporting period.



Application RTM Table — Shows response time statistics for all applications in the selected network group that experienced delay during the reporting period. You can re-sort the table by Application, Critical, Transactions, Good Trans, or TCP Conn by clicking on the corresponding column heading. This table contains the following information:

Statistic	Description
Application	The name of the application
Critical	Indicates whether the application is one that you have designated as critical .
Transactions	The total number of transactions reported for the application
Good Trans	The percentage of transactions that completed within the total delay threshold
TCP Conn Inits	The number of TCP connections established by the application and used by the transactions that were counted in the summary
Avg PET (ms)	The average Packet Exchange Time (PET) for transactions in the application. The PET represents the interval between a data packet leaving the PacketShaper and its acknowledgment (ACK) arriving.
Avg RTT (ms)	The average Round Trip Time (RTT) for packets in the application. The RTT represents the average number of milliseconds spent in transit when a client and server exchange the SYN (synchronize sequence numbers flag) and its corresponding ACK (acknowledge flag). A transaction involving a large amount of data requires the data to be divided into multiple packets. Whereas a transaction's network delay reflects the total transit time for all required packets, the RTT reflects the time for a single small packet to make its way from client to server and another packet to make the return trip. Use the RTT to determine if a large network delay is due to large transactions or a slow network. If the RTT is much smaller than the network delay, the transactions were large. If the two averages are close, a sluggish network caused the longer network delays.
Normalized Delay (ms)	The transaction delay in the network, normalized by transaction size. It shows how long it takes to send 1KB of data. This statistic allows an accurate comparison of response-time data for different applications or servers. Without normalizing the delay, response times vary depending on the size of the transaction. This statistic eliminates size as a factor of network delay
Avg Network Delay (ms)	The average response times in milliseconds of the traffic class over the reporting period. This statistic represents only the portion of the transaction time that is attributable to the network, enabling you to analyze network delay.
Avg Server Delay (ms)	The average response times in milliseconds of the traffic class over the reporting period. This graph shows only the portion of the transaction time that is attributable to the server, enabling you to analyze server delay.

Statistic	Description
Avg Total Delay (ms)	Average number of milliseconds to complete transactions; includes network delay and server delay.

Application	Critical	Transactions	Good Trans	TCP Conn	Avg PET (ms)	Avg RTT (ms)	Normalized Delay (ms)	Avg Network Delay (ms)	Avg Server Delay (ms)	Avg Total Delay (ms)
ProxySG: Management	No	84.26 K	100.00 %	85.12 K	11.19	2.51	533.92	35.78	11.89	47.67
HTTP	No	13.5 K	97.47 %	60.82 K	4.61	17.71	632.54	117.01	569.03	686.05
CIFS	Yes	6.62 K	89.89 %	1.49 K	12.28	0.93	2.21 K	40.28	5.31	45.59
FTP	Yes	3.4 K	99.94 %	4.9 K	137.89	139.76	8.77 K	922.91	7.44	930.34
SSL	No	2.55 K	99.96 %	46.45 K	36.52	66.59	3.81 K	108.34	34.84	143.16
DCOM	No	2.43 K	98.86 %	1.22 K	9.83	2.44	3.6 K	159.37	54.82	214.19
MS-Exchange	No	2.04 K	99.16 %	613	26.67	1.27	738.62	128.93	28.93	157.86
Symantec	No	1.51 K	7.17 %	683	1.96	1.10	1.83 K	89.12	45.56	134.68
Skype	No	1.31 K	100.00 %	118	33.84	91.02	84.23 K	5.85 K	5.4 K	11.25 K
SMTP	No	897	100.00 %	1.84 K	160.28	1.10	8.88 K	160.17	1.36	161.53
DNS	No	776	98.64 %	2.2 K	0.62	0.15	29.42	1.97	1.28	3.25
SOAP:HTTP	No	152	100.00 %	304	11.25	1.09	42.26	74.47	5.86	80.34
PostgreSQL	No	129	100.00 %	14	5.16	1.00	4.46	284.82	59.36	344.18
SSL-No-Cert	No	111	100.00 %	8	155.29	94.78	15.54 K	3.52 K	340.99	3.86 K
LDAP	No	110	97.27 %	110	19.75	23.61	82.20	128.87	1.32	130.19
FlashVideo	No	41	100.00 %	12	1.66	6.71	62.63	406.80	132.68	539.49
Facebook	No	19	100.00 %	38	26.47	10.89	30.05	226.21	8.87 K	9.09 K
MS-ActiveDir	No	9	100.00 %	16	89.50	19.56	208.95	272.11	1.67	273.78
Webshots	No	7	100.00 %	14	75.00	24.00	802.42	104.14	64.57	168.71
SMS	No	6	100.00 %	6	0.00	43.50	81.36	47.83	1.17	49.00
MPEG-Audio	No	5	100.00 %	2	1.92	7.00	0.87	321.40	78.40	399.80
ActiveX	No	4	100.00 %	4	0.00	28.25	30.96	28.25	1.00	29.25
WinMedia	No	1	100.00 %	3	4.00	0.00	1.15	2.00	4.00	6.00
Citrix	No	0	100.00 %	1	0.00	0.00	0.00	0.00	0.00	0.00
eDonkey	No	0	100.00 %	1	0.00	0.00	0.00	0.00	0.00	0.00
Lotus-IM	No	0	33.64 %	324	0.00	0.00	0.00	0.00	0.00	0.00
MSSQL	No	0	100.00 %	2	0.00	0.00	0.00	0.00	0.00	0.00
NetBIOS-IP	No	0	49.79 %	705	0.00	0.00	0.00	0.00	0.00	0.00
SSH	No	0	100.00 %	158	0.00	0.00	0.00	0.00	0.00	0.00
WebEx	No	0	100.00 %	2	0.00	0.00	0.00	0.00	0.00	0.00

Page: 1

Created on: Dec 2, 2009 3:23 PM

Click on an application name to drill down to the next level of detail (RTM statistics for each device).

Drill Down: Application Response Time for Application

When you drill down on an application from the top-level of the Application Response Time Report, a second-level report detailing application response on the devices that reported traffic for the selected application:


Top 25 Devices by Avg Total Delay Graph — Details the RTM statistics for each device that reported traffic for the selected application during the reporting period.

Application RTM per Device Table — Details the RTM statistics for each device that reported traffic for the selected application during the reporting period. For each device, the table shows the same RTM statistics that were displayed on the top-level view for the application as a whole.

Application Response Time Report for Application 'HTTP'

For View: 'Network'

From Sun, Nov 29, 2009 00:00:00 AM, GMT-0800 to Wed, Dec 2, 2009 03:00:00 PM, GMT-0800




IP Address	Device Name	Transactions	Good Trans	TCP Conn	Avg PET (ms)	Avg RTT (ms)	Normalized Delay (ms)	Avg Network Delay (ms)	Avg Server Delay (ms)	Avg Total Delay (ms)
10.9.50.91	France	11.37 K	94.19 %	25.63 K	7.95	15.18	654.94	117.26	662.69	779.95
10.9.59.70	East Coast	2.13 K	99.86 %	35.19 K	1.26	31.22	1.18 K	115.70	69.41	185.11



Click on a device entry to drill down to the next level of detail (application RTM for the selected device).





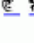




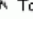
Drill Down: Application RTM for Device


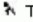
In the Device RTM report, you can click on a device entry to see the application RTM for the selected device.

Application RTM for Device Graphs — Shows two graphs that detail the RTM statistics on the selected device. The first graph compares the total delay, server delay, and network delay on the device. The second graph compares the round trip time (RTT) and the normalized network delay.



Application RTM for Device Table — Details RTM statistics for the traffic classes reported on the selected device. Click an icon to [drill-down](#) to top talkers  or listeners  for a class.

Traffic Class		Transactions	Good Trans	TCP Conn	Avg PET (ms)	Avg RTT (ms)	Normalized Delay (ms)	Avg Network Delay (ms)	Avg Server Delay (ms)	Avg Total Delay (ms)
/Outbound/HTTP	 	2.06 K	90.41 %	4.16 K	40.27	32.07	104.39	107.34	64.71	172.04
/Outbound	 	75	100.00 %	30.95 K	25.03	7.88	1.09 K	344.99	198.39	543.37
/NewYork/HTTP	 	0	100.00 %	75	0.30	0.00	0.00	0.00	0.00	0.00
/Inbound	 									
/NewYork/HTTP	 									

 To Top Listeners  To Top Talkers

Top Applications Report

Understanding what applications are running on your network is the first step to identifying and solving application delivery problems. The Top Applications report enables you to identify the top bandwidth-consuming applications that are running on your network so that you can ensure that your network resources are truly going to the mission critical applications that are central to your business.


When you configure the report, you specify how to determine the top applications — by number of bytes or packets — as well as the number of top applications to identify. This report provides several levels of detail.

Top-Level Views


The top level of this report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.

Pie Chart — Shows the bandwidth consumption of each top application relative to the others. Each pie slice represents the percentage of bandwidth the application consumed during the reporting period. Hover over a slice to see the associated application name and summary details. The legend to the right of the pie chart associates an application name with a pie slice color.

Application Utilization Table — For each of the top applications, this table shows the ranking of the application in bandwidth consumption, the name of the application, the percentage of total bandwidth consumed by the application, the total number of packets and bytes used by the application, the efficiency level percentage, and the average and peak rates for the application during the reporting period.





Average Rate Stack Chart — Shows throughput in the network for the top applications.



Two drill-down reports are available:

[Application Activity for a specific application](#) — Click on an application name to drill down to details about the specific devices and groups responsible for the application traffic.

[Top Talkers or Top Listeners for a specific application](#) — Click on an icon to display a list of top talkers  or top listeners  for the application. Note that you must be collecting FDR data in order to drill down to the top talkers or listeners report.

Top Immediate Children Report

The Top Immediate Children report allows you to identify the relative portions of bandwidth allocated to the ten most active children classes of a specified traffic class. This graph is similar to the [Top Traffic Classes Summary Report](#) except that this report displays direct children only, while the Top Traffic Classes Summary Report shows leaf classes (classes that don't have any children of their own). In other words, the Top Immediate Children report allows you to graph a traffic class' children, without grandchildren.


When you configure the report, you specify how to determine the top child classes—by number of bytes or packets — as well as the number of top child classes to display. In addition, you must specify the parent class on which you want to report and the specific [network group or view](#) to which to restrict the report.



Top-Level Views

The top level of this report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report. This report contains the following elements:

Pie Chart — Shows each child class' average bandwidth usage in bits per second and its percentage of the parent class' total bandwidth usage. Each pie slice represents the percentage of bandwidth the traffic class consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates a traffic class with a pie slice color.

Child Class Utilization Table — For each of the top children, this table shows the ranking of the class in bandwidth consumption, the name of the class, the percentage of the parent class' total bandwidth usage, the total number of packets and bytes used by the class, the partition size and utilization, the efficiency percentage (ratio of bytes not requiring retransmission to the total number of bytes sent), and the average and peak rates for the class during the reporting period.



Clicking on the icon allows you to [drill down](#) to a list of Top Talkers  or Top Listeners  for the class. You can also click on a traffic class name to [drill down](#) into device details. This drill down is based on FDR data and you will therefore only be able to drill down if you are collecting FDR. In addition, because flow details are only reported for leaf classes, you will only be able to drill down on a child class if it is a leaf class.

Device Reports

Device Compression Summary Report

For an overview of how compression is working on a device or group of devices, you can [run this report](#). This report displays compression statistics for the Inbound and Outbound links on the devices in the [network branch](#) or [view](#) that you specify when you configure the report. For specific information on how to solve compression problems, refer to [Compression Troubleshooting](#) in PacketGuide.

This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.

Top-Level Views

The top level of this report has a table that summarizes compression statistics for the Inbound and Outbound links for the devices in the selected network branch or view. The Inbound statistics represent inbound traffic that gets decompressed by Xpress, and the Outbound statistics represent outbound traffic that Xpress compresses. This table contains the following information:

Statistic	Description
Device	The IP address of the device. Note that there are two entries for each device; one for inbound statistics and one for outbound statistics.
Direction	The link direction (Inbound or Outbound)
Comp	Off indicates that Xpress compression is disabled; On indicates that Xpress compression is enabled
Bytes Saved (%)	The percentage of bytes saved on the link due to compression
Postcomp	Postcompression bytes — the number of bytes that went through a compression tunnel
Precomp	Precompression bytes — the number of bytes that would have passed through the link if compression wasn't enabled
Bytes Saved	The number of bytes that didn't have to traverse the link, due to compression; it's the difference between precompression and postcompression bytes.
Non-Comp	Noncompressible bytes — the number of bytes that PacketWise did not attempt to compress, either because they didn't belong to a compressible service or because they were destined for a location without an Xpress partner.
Post Avg	Postcompression average rate — the throughput rate for compressed

Statistic	Description
Rate	data in bits per second
Pre Avg Rate	Precompression average rate — the throughput rate (in bits per second) that would have been required if compression was not enabled
Post Peak Rate	Postcompression peak rate — the highest throughput rate (in bits per second) used for the data after it was compressed
Pre Peak Rate	Precompression peak rate — the highest throughput rate (in bits per second) that would have been required if compression was not enabled


Device Compression Summary Report											
For View: 'Network'											
From Fri., Oct 01, 2010 12:00:00 AM, GMT-0700 to Tue., Oct 12, 2010 12:00:00 AM, GMT-0700											
Device	Direction	Comp	Bytes Saved [%]	Postcomp	Precomp	Bytes Saved	Non-Comp	Post Avg Rate	Pre Avg Rate	Post Peak Rate	Pre Peak Rate
10.78.52.51	Inbound	On	0.00 %	706.03 G	706.03 G	0	706.03 G	76.43 M	76.43 M	224.17 M	224.17 M
10.78.52.51	Inbound	On	0.00 %	7.15 T	7.15 T	0	7.15 T	165.43 M	165.43 M	406.66 M	406.66 M
10.78.52.51	Outbound	On	0.00 %	404.75 G	404.75 G	0	404.75 G	43.81 M	43.81 M	99.15 M	99.15 M
10.78.52.51	Outbound	On	0.00 %	3.94 T	3.94 T	0	3.94 T	91.28 M	91.28 M	243.27 M	243.27 M
10.78.52.50	Inbound	On	0.00 %	340.33 G	340.33 G	0	340.33 G	79.62 M	79.62 M	386.26 M	386.26 M
10.78.52.50	Inbound	On	0.00 %	7.22 T	7.22 T	0	7.22 T	167.17 M	167.17 M	416.8 M	416.8 M
10.78.52.50	Outbound	On	0.00 %	209.3 G	209.3 G	0	209.3 G	48.97 M	48.97 M	386.46 M	386.46 M
10.78.52.50	Outbound	On	0.00 %	4.07 T	4.07 T	0	4.07 T	94.27 M	94.27 M	247.35 M	247.35 M
10.78.52.49	Inbound	On	0.00 %	777.79 G	777.79 G	0	777.79 G	74.36 M	74.36 M	386.46 M	386.46 M
10.78.52.49	Inbound	On	0.00 %	7.23 T	7.23 T	0	7.23 T	167.43 M	167.43 M	416.68 M	416.68 M
10.78.52.49	Outbound	On	0.00 %	463.68 G	463.68 G	0	463.68 G	44.33 M	44.33 M	368.21 M	368.21 M
10.78.52.49	Outbound	On	0.00 %	4.11 T	4.11 T	0	4.11 T	95.16 M	95.16 M	247.25 M	247.25 M
Page: 1	Created on: Oct 12, 2010 5:28 PM										

Clicking on a device address allows you to drill down to graphs that detail how compression is working on the device.


Drill Down: Compression Graphs

When you drill down on a particular device in the Device Compression Summary, the following graphs and tables are displayed:


Inbound Average Rate Graph — Compares average bandwidth usage of compressible traffic on the inbound link of the selected device, with and without compression. The *Postcompression Avg Rate* line represents usage with compression enabled and the *Precompression Avg Rate* line represents what average usage would have been without compression.




Outbound Average Rate Graph — Compares average bandwidth usage of compressible traffic on the outbound link of the selected device, with and without compression. The *Postcompression Avg Rate* line represents usage with compression enabled and the *Precompression Avg Rate* line represents what average usage would have been without compression.



Inbound Peak Rate Graph — Compares peak bandwidth usage of compressible traffic on the inbound link of the selected device, with and without compression. The *Postcompression Avg Rate* line represents usage with compression enabled and the *Precompression Avg Rate* line represents what average usage would have been without compression.



Outbound Peak Rate Graph — Compares peak bandwidth usage of compressible traffic on the outbound link of the selected device, with and without compression. The *Postcompression Avg Rate* line represents usage with compression enabled and the *Precompression Avg Rate* line represents what average usage would have been without compression.



The graphs on the second and third pages compare the top compressible traffic classes. To move between pages, click the blue arrow icons at the top of the report.

Top 10 Inbound Compressible Traffic Classes Bar Graph and Table — The bar graph shows the percentage of bytes saved for each of the top 10 compressible inbound traffic classes, due to compression. The table ranks each traffic class in order of greatest percentage saved. The percentage of bytes saved value is calculated by subtracting precompression bytes (the size without any compression) and postcompression bytes (the size after compression) and dividing this difference by precompression bytes. For example, if a class' traffic is 400k without compression and is compressed to 100k, the Compression Percent Bytes Saved would be 75%. Note that only compressible traffic is considered in these calculations.

Top 10 Outbound Compressible Traffic Classes Bar Graph and Table — The bar graph shows the percentage of bytes saved for each of the top 10 compressible outbound traffic classes, due to compression. The table ranks each traffic class in order of greatest percentage saved. The percentage of bytes saved value is calculated by subtracting precompression bytes (the size without any compression) and postcompression bytes (the size after compression) and dividing this difference by precompression bytes. For example, if a class' traffic is 400k without compression and is compressed to 100k, the *Compression Percent Bytes Saved* would be 75%. Note that only compressible traffic is considered in these calculations.

Link Utilization Report

The Link Utilization report provides a quick way for you to analyze link utilization on your inbound and/or outbound links to help you troubleshoot network issues and plan for future bandwidth needs. There are a couple of ways you can use this report:

- **Troubleshooting of current issues**—If your users are reporting network delays and time-outs, you might want to run this report to try and determine the cause of the issue. For example, drops in network efficiency indicate high retransmission rates, which is a sign of congested router queues. By drilling-down into the data, you can identify which PacketShaper is reporting the link issues. From there you could run top talker or top traffic class reports for that device during the same time range to help you determine if the problem is a bursty traffic class that is monopolizing the link or if the link issues are instead caused by a router or other network device.
- **Capacity planning for the future**—If you enable trending when you run this report, IC uses the link utilization history to forecast future usage so that you can estimate if and when you will outgrow your WAN capacity. If the trending report shows that you are continually using more and more of your WAN capacity you can get an idea of when need to add additional bandwidth or change your traffic management strategy. For example, perhaps you've been overly generous in allowing employees unrestricted access to Internet radio, but if that means an expensive bandwidth upgrade is imminent, then you might want to restrict access instead.


When you configure the report, you specify whether or not to include the trending data to help you with capacity planning.

This report displays line graphs and tables that detail inbound and/or outbound link utilization percentages, peak and average bandwidth usage, and network efficiency values for the selected network group, subgroup, device, or view during the reporting period. If you enabled trending, this report also provides link utilization forecast data to help you with future capacity planning. The graphs show aggregated data for all devices in the reporting group, whereas the tables show a breakdown of the link utilization statistics by PacketShaper. If you chose to include both inbound and outbound data in the report, there are separate tables for each direction. This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.

This report includes the following graphs and tables. Note that the graphs will show data for inbound, outbound, or both inbound and outbound data, depending on how you configured the report. In this example, all graphs show data for both inbound and outbound links. In addition, in this example Trending is enabled.


Link Utilization Forecast

If you enabled Trending and your reporting period is greater than one day, the report shows the Link Utilization Forecast. This graph shows the expected link utilization percentage over the specified time range based on historical link utilization data. Use this graph to help you plan future capacity needs. In this example, link utilization is projected to stay steady. If you notice your link utilization forecast trending up to your capacity limits, you may need to either change your policies to be more strict about bandwidth usage or you should start planning for a bandwidth upgrade.




Link Utilization

The Link Utilization graph shows the percentage of your total link capacity that is being used. Use this graph to determine if and when you are bumping up against your capacity limit.




Average Rate

Shows the average bandwidth usage in bits per second for all managed links in the selected reporting group.



Peak Rate

Shows the peak bandwidth usage in bits per second for all managed links in the selected reporting group. Peak rates tend to occur when applications and protocols *burst* — rapidly expand to consume great quantities of bandwidth, undermining the performance of applications that are, perhaps, more urgent. This behavior doesn't necessarily imply unruly users or bad applications. The applications might be critical but tend to consume greater than an appropriate share of bandwidth (Microsoft Exchange, for example). Or they might be important, but not interactive and not particularly urgent (FTP, for example). Or they might be recreational applications (music downloads, for example).




Note: Because ME variables are collected in sample intervals, they are meant to be examined over time rather than instantaneously. Because of this you may sometimes see cases where the peak values that are shown are less than the average values. This is usually occurs when you have a peak rate of zero and a small average rate that appear immediately after an interval with a large amount of traffic. This is because average rate is a weighted quantity, meaning that each second of stored data includes exponentially weighted values from the previous five seconds. Therefore, if you have a traffic interval when no data passed through (peak rate of zero), but with an average rate that includes a small weighted number from the last few seconds of the previous

one-minute interval, when graphed, will show an average rate in excess of the peak rate. Although this may seem like it is incorrect, it is just really a factor how ME data is collected.

Network Efficiency

Shows the percentage of all bytes that are *not* TCP retransmits, giving you an indication of overall network health. For example, a value of 95% indicates that of all traffic on the network (TCP and UDP), 5% was retransmitted TCP traffic. (UDP is assumed to be 100% efficient since retransmissions are not calculated on this type of traffic.) Retransmissions, traffic that must traverse the network multiple times for successful arrival, should optimally be as close to zero as possible. But when router queues deepen and cause dropped packets, retransmissions spike. When latency increases the frequency of time-outs, retransmissions spike. When a busy IP network behaves precisely as designed under heavy loads, retransmissions spike.

Check to see whether drops in network efficiency coincide with peak rates. If so, this could indicate that you might have some greedy applications that are spiking periodically, slowing more critical applications. Dips in network efficiency may also indicate router congestion, causing delays and time-outs.



Inbound/Outbound Link Utilization Tables

The link utilization tables show a breakdown of the link utilization statistics by PacketShaper. For each PacketShaper, the tables display the link size and the percentage of peak and average link usage, the network efficiency percentage, and the peak and average bandwidth usage. You can re-sort the table data by clicking on **Shaper IP**, **Shaper Name**, **Link Size** or **Link Utilization** in the table header. Depending on how you configured the report, it will display either the Inbound Link Utilization Table, the Outbound Link Utilization Table, or both tables (as in this example). If your report includes data from more than one PacketShaper, you can also drill-down on an IP address to view graphs that detail link utilization on that particular appliance.

Inbound Link Utilization Table							
Shaper IP	Shaper Name	Link Size	Link Utilization	Peak Link utilization	Network Efficiency	Avg Rate (bps)	Peak Rate (bps)
200.200.221.12	PS12	100.0 M	2.90 %	6.05 %	97.08 %	2.9 M	6.05 M
200.200.221.14	PS-14	100.0 M	1.96 %	3.18 %	95.94 %	1.96 M	3.18 M
200.200.221.13	PS-13	100.0 M	2.97 %	6.20 %	97.18 %	2.97 M	6.2 M
Outbound Link Utilization Table							
Shaper IP	Shaper Name	Link Size	Link Utilization	Peak Link utilization	Network Efficiency	Avg Rate (bps)	Peak Rate (bps)
200.200.221.12	PS12	100.0 M	4.56 %	16.77 %	99.02 %	4.56 M	16.77 M
200.200.221.14	PS-14	100.0 M	2.65 %	6.83 %	98.77 %	2.65 M	6.83 M
200.200.221.13	PS-13	100.0 M	4.73 %	17.67 %	99.06 %	4.73 M	17.67 M

TCP Health Report


The TCP Health report gives you a comprehensive picture of TCP connections for a particular network group, sub-group, device, or view. It compares the number of TCP connections that were started, aborted, ignored by the server (that is, the server never responded), and refused by the server during the specified reporting period.

For example, suppose you [created network views](#) for each branch office, and one of your offices complained about a server problem (such as slow speed, or many disconnects, or hung systems). By comparing TCP Health reports for the various branch offices, you can see what the normal ranges are for *Connections* and *Server Ignores* on your network and use this information to figure out what is causing the server problem at the one branch office.

When you configure the report, you specify a network branch or view and time period for which to monitor TCP connections. This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.

Top-Level Views

TCP Health Graph — Shows the overall performance of the traffic to and from the WAN or Internet connection. Individual lines in the graph represent the total number of TCP connections, TCP aborted connections, TCP attempted connections that were ignored by the server, and TCP connections that were refused by the server.



TCP Health Table — For each of the traffic classes in the selected network branch or view that generated TCP connections during the reporting period, this table shows the following statistics:

Statistic	Description
Traffic Class	The name of the traffic class
Bytes	The total number of bytes transmitted by the traffic class during the reporting period
Packets	The total number of packets transmitted by the traffic class during the reporting period
New Connections	The number of TCP connections achieved by the traffic class during the reporting period
Aborts	The number of TCP connections that were aborted, rather than closed, during the reporting period
Server Ignores	The number of times a server ignored a request for a TCP connection during the reporting period. Large and sudden increases in this statistic may indicate a network attack, such as a Denial of Service (DoS) attack, or may indicate other network problems such as an overloaded server.
Server Refuses	The total number of times during the reporting period that a server refused a TCP connection request from the traffic class. Large and sudden increases in this statistic may indicate unusual network activity, such as a port scan or DoS attack.

Traffic Class	Bytes	Packets	New Connections	Aborts	Server Ignores	Server Refuses
Inbound	252.8 G	267.44 M	9.47 M	120.42 K	91.58 K	1.29 K
Inbound/Default	240.68 G	242.06 M	8.56 M	114.25 K	84.32 K	664
Inbound/CIFS	12.07 G	18.48 M	10.6 K	123	7.03 K	180
Outbound	3.53 G	9.76 M	6.85 M	120.28 K	91.52 K	1.29 K
Outbound/Localhost	1.91 G	4.04 M	666.13 K	5.26 K	90	450
Outbound/Default	1.42 G	4.03 M	6.16 M	114.51 K	80.19 K	658
Outbound/Localhost/FlowRecords	960.03 M	634.1 K	0	0	0	0
Inbound/Localhost	924.09 M	4.64 M	862.14 K	5.26 K	74	450
Outbound/Localhost/HTTP	877.03 M	3.68 M	555.16 K	1.3 K	90	27
Inbound/Localhost/HTTP	751.69 M	4.0 M	751.46 K	1.3 K	74	27
Inbound/Categories	220.88 M	224.71 K	18.66 K	353	0	0
Inbound/DNS	203.73 M	1.17 M	4.25 K	6	49	0
Inbound/Skype	110.38 M	105.23 K	817	15	0	0
Inbound/Localhost/FTP	100.65 M	176.59 K	26	15	0	0
Inbound/mDNS	88.82 M	617.34 K	0	0	0	0
Inbound/OSPF	74.24 M	951.81 K	0	0	0	0
Outbound/Localhost/SSL	50.29 M	192.89 K	23.57 K	208	0	0
Outbound/CIFS	43.76 M	99.04 K	3.97 K	121	7.01 K	180
Inbound/Categories/Reference	43.46 M	59.8 K	1.91 K	12	0	0
Inbound/Localhost/DNS	35.46 M	227.08 K	0	0	0	0
Inbound/Localhost/SSL	31.92 M	210.88 K	23.55 K	208	0	0
Outbound/DNS	31.52 M	489.63 K	4.26 K	6	49	0

Click on a traffic class name to drill down to details about the specific devices that reported the TCP traffic in the selected class.

Drill Down: TCP Health by Device

On the TCP Health report, you can click on a traffic class name to drill down to details about the specific devices that reported the TCP traffic in the selected class.

Devices for Class — Details the TCP health statistics for each device that reported TCP statistics for the selected traffic class during the reporting period. This table shows the following statistics for each reporting device:

Statistic	Description
IP address	The IP address of the device
Device	The host name of the device
Bytes	The total number of bytes transmitted by the device during the reporting period
Packets	The total number of packets transmitted by the device during the reporting period
New Connections	The number of TCP connections achieved by the device during the reporting period
Aborts	The number of TCP connections that were aborted, rather than closed, during the reporting period
Server Ignores	The number of times a server ignored a request for a TCP connection during the reporting period. Large and sudden increases in this statistic may indicate a network attack, such as a Denial of Service (DoS) attack, or may indicate other network problems such as an overloaded server.
Server Refuses	The total number of times during the reporting period that a server refused a TCP connection request from the device. Large and sudden increases in this statistic may indicate unusual network activity, such as a port scan or DoS attack.


TCP Health Report							
For View: 'Network'							
For Traffic Class: '/Inbound/Categories/Computers_Internet'							
From Sun., Aug 1, 2010 00:00:00 AM, GMT-0700 to Wed., Sep 1, 2010 00:00:00 AM, GMT-0700							
IP Address	Device Name	Bytes	Packets	New Connections	Aborts	Server Ignores	Server Refuses
10.9.59.70	East Coast	30.25 M	22.27 K	1.51 K	96	0	0
Page: 1		Created on: Sep 13, 2010 2:38 PM					

Click on a device entry to drill down to a graph that illustrates TCP health on the device.

Drill Down: Device TCP Health

On the TCP Health by Device report, you can click on a device entry to drill down to a graph that illustrates TCP health on the device.

Connections for Device — Shows the overall performance of the TCP traffic on the selected device. Individual lines in the graph represent the total number of new TCP connections, TCP aborted connections, TCP attempted connections that were ignored by the server, and TCP connections that were refused by the server.



Top Traffic Classes Summary Report

Knowing the identity of traffic running over your network is a big first step in managing and controlling the performance of network applications. The Top Traffic Classes Summary report enables you to identify the top bandwidth-consuming traffic classes that are running on your network.


When you configure the report, you specify how to determine the top traffic classes — by number of bytes or packets — as well as the number of top traffic classes to identify. This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.

Top-Level Views


Pie Chart — Shows the bandwidth consumption of each top traffic class relative to the others. Each pie slice represents the percentage of bandwidth the traffic class consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates a traffic class name with a pie slice color.

Traffic Class Utilization Table — For each of the top traffic classes, this table shows the ranking of the traffic class in bandwidth consumption, the name of the traffic class, the percentage of total bandwidth consumed by the class, the total number of packets and bytes used by the class, the efficiency level percentage, and the average and peak rates for the traffic class during the reporting period. Clicking on a traffic class name allows you to [drill down](#) to details about the specific devices and groups responsible for generating the traffic that is classified into the

selected traffic class. Clicking on the icon allows you to [drill down](#) to a list of Top Talkers or Top Listeners for the class. This drill down is based on FDR data and you will therefore only be able to drill down if you are collecting FDR. In addition, because flow details are only reported for leaf classes, you will only be able to drill down on a child class if it is a leaf class.



Average Rate Stack Chart — Shows throughput in the network for the top traffic classes.



Traffic Class Compression Report


For an overview of how compression is working on a network group or view, you can [run this report](#). This report displays compression statistics all of the views and traffic classes in the network group, sub-group, device, or view that you specify when you configure the report. This report allows you to drill down and see how compression is working on specific devices and on specific traffic classes on a device. It is useful for helping you troubleshoot compression problems by enabling you to identify the troublesome hosts and/or applications. For specific information on how to solve compression problems, refer to [Compression Troubleshooting](#) in PacketGuide.

This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.


Top-Level Views

The top level of this report shows graphs and a table that summarizes compression statistics for the network groups and traffic classes that belong to the selected network group or view as follows:


Average Rate Graph — Compares average bandwidth usage of compressible traffic on the of the selected network branch or view, with and without compression. The *Postcompression Avg Rate* line represents usage with compression enabled and the *Precompression Avg Rate* line represents what average usage would have been without compression.



Bytes Saved By Compression Graph — Compares bytes saved with and without compression



Bytes Saved By Compression Percent Graph — Compares percentages of bytes saved for compressible and total bytes




Traffic Class Compression Table — The table contains the following information:

Statistic	Description
Group/Traffic Class	The name of the group or traffic class for which compression statistics are displayed. Note that group statistics are shown on the top table and traffic class statistics are shown on the bottom the table.
Postcomp Bytes	The number of bytes that went through a compression tunnel
Precomp Bytes	The number of bytes that would have passed through the link if compression wasn't enabled
Bytes Saved	The number of bytes that didn't have to traverse the link, due to compression; it's the difference between precompression and postcompression bytes.
Bytes Saved (%)	Bytes saved by compression as a percentage of total bytes

Statistic	Description
Compressible Bytes	The total number of compressible bytes in the group or traffic class
Comp Bytes Saved	Bytes saved by compression as a percentage of total compressible bytes

Traffic Class Compression						
Showing page 1 of 38						
Group	Postcomp Bytes	Precomp Bytes	Bytes Saved	Bytes Saved (%)	Compressible Bytes	Comp Bytes Saved
Network	2.25 T	2.25 T	0	0.00 %	0	0.00
Traffic Class	Postcomp Bytes	Precomp Bytes	Bytes Saved	Bytes Saved (%)	Compressible Bytes	Comp Bytes Saved
/Outbound/medium/rexec	0	0	0	0.00 %	0	0.0
/Outbound/medium	0	0	0	0.00 %	0	0.0
/TVU-Player						
/Outbound/medium	0	0	0	0.00 %	0	0.0
/Symantec						
/Outbound/medium/Skype	0	0	0	0.00 %	0	0.0
/Outbound/medium	0	0	0	0.00 %	0	0.0
/SSL-No-Cert						
/Outbound/medium/SSL	5.89 M	5.89 M	0	0.00 %	0	0.0
/Outbound/medium/SMTP	163.47 K	163.47 K	0	0.00 %	0	0.0
/Outbound/medium/RTP-I	0	0	0	0.00 %	0	0.0
/RTP-HG729						
/Outbound/medium/RTP-I	644.31 K	644.31 K	0	0.00 %	0	0.0
/RTP-I-Dynamic						
/Outbound/medium/RTP-I	0	0	0	0.00 %	0	0.0
/Default						
/Outbound/medium/RTP-I	644.31 K	644.31 K	0	0.00 %	0	0.0
/Outbound/medium/POP3	0	0	0	0.00 %	0	0.0
/Outbound/medium	0	0	0	0.00 %	0	0.0
/PEPGate						
/Outbound/medium	222.77 K	222.77 K	0	0.00 %	0	0.0
/NetBIOS-IP						
/Outbound/medium/NFS	0	0	0	0.00 %	0	0.0
/Outbound/medium/MSSQL	0	0	0	0.00 %	0	0.0
/Outbound/medium	0	0	0	0.00 %	0	0.0
/MS-Exchange						

Click the blue arrow  icon to move to the next page of data. To drill-down to the next level of detail, click on a group or traffic class name in the table.

Drill Down: Traffic Class Compression by Device

Devices for Traffic Class — This view shows compression details for the specific devices in the selected group or traffic class. See the table above for information about a specific field.

Traffic Class Compression Report - Traffic Classes for Device '10.78.53.135'						
From Thu., Sep 16, 2010 12:00:00 AM GMT-0700 to Thu., Sep 16, 2010 03:00:00 PM GMT-0700						
Traffic Class	Postcomp Bytes	Precomp Bytes	Bytes Saved	Bytes Saved (%)	Compressible Bytes	Comp Bytes Saved
10.78.53.135/Outbound/medium/ICMP	110.61 M	110.61 M	0	0.0 %	0	0.0
Child Traffic Class	Postcomp Bytes	Precomp Bytes	Bytes Saved	Bytes Saved (%)	Compressible Bytes	Comp Bytes Saved
There are no records for this table.						
Page: 1 Created on: Sep 16, 2010 3:08 PM						

To drill down to compression information for the specific traffic classes for a device, click on the IP address for the device you are interested in.

Drill Down: Traffic Class Compression for Device

Traffic Classes for Device — This view shows compression details for the specific traffic classes on the selected device. See the table above for information about a specific field.

Traffic Class Compression Report - Traffic Classes for Device '10.78.53.135'						
From Thu., Sep 16, 2010 12:00:00 AM GMT-0700 to Thu., Sep 16, 2010 03:00:00 PM GMT-0700						
Traffic Class	Postcomp Bytes	Precomp Bytes	Bytes Saved	Bytes Saved (%)	Compressible Bytes	Comp Bytes Saved
10.78.53.135/Outbound /medium/ICMP	110.61 M	110.61 M	0	0.0 %	0	0.0
Child Traffic Class	Postcomp Bytes	Precomp Bytes	Bytes Saved	Bytes Saved (%)	Compressible Bytes	Comp Bytes Saved
There are no records for this table.						
Page: 1 Created on: Sep 16, 2010 3:08 PM						

Traffic Class Response Time Report

Response time measurement (RTM) provides information about the amount of time connection-based TCP traffic spends traveling between a client and a server and the time used by the server itself. This allows you to investigate response times and identify the source of network delays. The Traffic Class RTM report contrasts RTM statistics in each traffic class. When you configure the report, you specify the network group, device, or view to which to restrict the report. For example, you could restrict the report so that it shows application utilization at a specific branch office or within a specific department.

This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.


Top-Level View

Traffic Class RTM Table — Displays a table that shows response time statistics for all traffic classes that have been discovered by IC. This table contains the following information:

Statistic	Description
Traffic Class	The name of a specific traffic class discovered by IntelligenceCenter on one of the reporting devices
Transactions	The total number of transactions reported for the traffic class
Good Trans	The percentage of transactions that completed within the total delay threshold
TCP Conn	The number of TCP connections established by the class and used by the transactions that were counted in the summary
Avg PET (ms)	The average Packet Exchange Time (PET) for transactions in the class. The PET represents the interval between a data packet leaving the PacketShaper and its acknowledgment (ACK) arriving.
Avg RTT (ms)	The average Round Trip Time (RTT) for packets in the class. The RTT represents the average number of milliseconds spent in transit when a client and server exchange the SYN (synchronize sequence numbers flag) and its corresponding ACK (acknowledge flag). A transaction involving a large amount of data requires the data to be divided into multiple packets. Whereas a transaction's network delay reflects the total transit time for all required packets, the RTT reflects the time for a single small packet to make its way from client to server and another packet to make the return trip. Use the RTT to determine if a large network delay is due to large transactions or a slow network. If the RTT is much smaller than the network delay, the transactions were large. If the two averages are close, a sluggish network caused the longer network delays.

Normalized Delay (ms)	The transaction delay in the network, normalized by transaction size. It shows how long it takes to send 1KB of data. This statistic allows an accurate comparison of response-time data for different applications or servers. Without normalizing the delay, response times vary depending on the size of the transaction. This statistic eliminates size as a factor of network delay
Avg Network Delay (ms)	The average response times in milliseconds of the traffic class over the reporting period. This statistic represents only the portion of the transaction time that is attributable to the network, enabling you to analyze network delay.
Avg Server Delay (ms)	The average response times in milliseconds of the traffic class over the reporting period. This graph shows only the portion of the transaction time that is attributable to the server, enabling you to analyze server delay.
Avg Total Delay (ms)	Average number of milliseconds to complete transactions; includes network delay and server delay

Traffic Class Response Time Report									
For View: 'Network'									
From Thu., Sep 16, 2010 03:00:00 PM, GMT-0700 to Thu., Sep 16, 2010 04:00:00 PM, GMT-0700									
Traffic Class	Transactions	Good Trans	TCP Conn	Avg PET (ms)	Avg RTT (ms)	Normalized Delay (ms)	Avg Network Delay (ms)	Avg Server Delay (ms)	Avg Total Delay (ms)
/Inbound/medium/SSL	0	100.00 %	17	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/medium/SMTP	0	75.00 %	4	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/medium/HTTP	0	99.95 %	1.87 K	17.52	0.00	0.00	0.00	0.00	0.00
/Inbound/medium/Default	0	0.00 %	4	15.00	0.00	0.00	0.00	0.00	0.00
/Inbound/medium/CIFS	0	100.00 %	1	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/medium	0	99.58 %	1.89 K	17.52	0.00	0.00	0.00	0.00	0.00
/Inbound/low/SSL	0	65.93 %	91	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/low/SMTP	0	33.33 %	15	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/low/MSN-Messenger	0	100.00 %	39	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/low/HTTP	0	91.99 %	674	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/low/FTP	0	100.00 %	1	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/low/Default	0	76.92 %	39	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/low/CIFS	0	100.00 %	1	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/low	0	87.91 %	860	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/high/SSL	0	98.39 %	1.66 K	36.78	0.00	0.00	0.00	0.00	0.00
/Inbound/high/SMTP	0	55.45 %	220	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/high/MSSQL	0	100.00 %	1	0.00	0.00	0.00	0.00	0.00	0.00
/Inbound/high/HTTP	0	96.86 %	1.47 K	49.48	0.00	0.00	0.00	0.00	0.00
/Inbound/high/Default	1.01 K	96.88 %	480	16.11	19.18	238.15	129.59	7.27	136.86

Click the blue arrow  icon to move to the next page of data. Click on a traffic class name to drill down to the next level of detail.

Drill Down: Traffic Class RTM per Device

In the Traffic Class RTM report, you can click on a class name to see RTM statistics for each device.

Traffic Class RTM per Device — Details the RTM statistics for each PacketShaper that reported traffic for the selected traffic class during the reporting period. For each device, the table shows the same RTM statistics that were displayed on the top-level view for the traffic class as a whole.


Traffic Class Response Time Report for Traffic Class 'Inbound/chico'										
For View: 'Network'										
From Thu., Sep 16, 2010 03:00:00 PM, GMT-0700 to Thu., Sep 16, 2010 04:00:00 PM, GMT-0700										
IP Address	Device Name	Transactions	Good Trans	TCP Conn	Avg PET (ms)	Avg RTT (ms)	Normalized Delay (ms)	Avg Network Delay (ms)	Avg Server Delay (ms)	Avg Total Delay (ms)
10.78.53.182	Lab-PS10k	4.54 K	0.00 %	17.68 K	948.16	227.54	2.21 K	2.87 K	3.45 K	6.32 K
10.78.53.181	Lab-PS10K-2	5.38 K	0.00 %	17.98 K	954.87	228.06	1.5 K	3.3 K	3.38 K	6.68 K

Page: 1 Created on: Sep 16, 2010 4:39 PM

Click on a device entry to drill down to the next level of detail (traffic class RTM for the selected device).

Drill Down: Traffic Class RTM for Device

Traffic Class RTM for Device Graph — Shows two graphs that detail the RTM statistics on the selected device. The first graph compares the total delay, server delay, and network delay on the device. The second graph compares the RTT and the normalized network delay.



Traffic Class RTM for Device Table — Details the RTM statistics for each child traffic class that reported traffic for the selected device and traffic class during the reporting period. For each child class, the table shows the same RTM statistics that were displayed on the top-level view for the traffic class as a whole

Child Traffic Class	Transactions	Good Trans	TCP Conn	Avg PET (ms)	Avg RTT (ms)	Normalized Delay (ms)	Avg Network Delay (ms)	Avg Server Delay (ms)	Avg Total Delay (ms)
/Inbound/chico	0	93.39 %	121	789.05	0.00	0.00	0.00	0.00	0.00
/mySpace									
/Inbound/chico	527	71.30 %	5.07 K	3.91 K	766.98	19.6 K	4.7 K	1.26 K	5.96 K
/eDonkey									
/Inbound/chico/SSL	0	62.62 %	1.08 K	54.83	0.00	0.00	0.00	0.00	0.00
/Inbound	0	0.00 %	178	868.42	0.00	0.00	0.00	0.00	0.00
/chico/SMTP									
/Inbound/chico	6	56.73 %	208	272.72	1.41 K	92.86 K	37.05 K	65.00	37.11 K
/KaZaA									
/Inbound	18	37.37 %	6.4 K	164.93	209.17	82.16	596.56	874.06	1.47 K
/chico/HTTP									
/Inbound/chico	2	0.00 %	207	1.19 K	451.50	2.06 K	1.98 K	803.50	2.78 K
/Gnutella									
/Inbound/chico/FTP	0	0.00 %	240	672.05	0.00	0.00	0.00	0.00	0.00
/Inbound/chico	4.81 K	0.00 %	4.4 K	1.34 K	165.98	1.3 K	3.1 K	3.6 K	6.71 K
/Delebit									
/Inbound/chico	17	49.32 %	73	1.1 K	666.41	1.05 K	5.1 K	10.22 K	15.33 K
/BitTorrent									

Page: 1 Created on: Sep 16, 2010 4:41 PM


Traffic Class Utilization Report

The Traffic Class Utilization report provides a quick way to get breakdown of the top-bandwidth consuming child traffic classes associated with a particular traffic class on a selected device, sub-group, or group over the reporting period. This report displays line graphs and tables that detail peak and average bandwidth usage and network efficiency values for the selected traffic class within the selected network group, subgroup, device, or view during the reporting period. The graphs show aggregated data for all devices in the reporting group, whereas the tables show a breakdown of the traffic class utilization statistics by PacketShaper and by group/subgroup. This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.

This report includes the following graphs and tables.


Average Rate vs. Peak Rate

Shows the selected traffic classes average and peak bandwidth consumption, in bits per second, over the reporting period.



Network Efficiency

Shows the percentage of all bytes that are *not* TCP retransmits, giving you an indication of how well the selected traffic classes and its child classes are performing. For example, a value of 95% indicates that of all traffic on the network (TCP and UDP), 5% was retransmitted TCP traffic. (UDP is assumed to be 100% efficient since retransmissions are not calculated on this type of traffic.) Retransmissions, traffic that must traverse the network multiple times for successful arrival, should optimally be as close to zero as possible.



Class Utilization Tables

The class utilization tables show a breakdown of the traffic class utilization statistics by PacketShaper and by group/sub-group. The Traffic Class utilization tables display a list of the devices and/or groups/sub-groups that generated traffic for the specified class and network group during the reporting period you configured. For each device or group entry, the table displays the total number of packets and bytes used by the device or group, the partition size and percentage of utilization, the efficiency level percentage, and the average and peak rates for the traffic generated by the device or group.

Class Utilization Device Table								
IP Address	Device Name	Bytes	Packets	Partition Size	Partition Util (%)	Efficiency	Avg Rate (bps)	Peak Rate (bps)
10.125.32.20	117-10014614	787.29 M	1.7 M	1.5 M	4.87 %	99.90 %	73.03 K	35.97 M
Class Utilization Group Table								
Group	Bytes	Packets	Partition Size	Partition Util (%)	Efficiency	Avg Rate (bps)	Peak Rate (bps)	
Network	787.29 M	1.7 M	1.5 M	4.87 %	99.90 %	73.04 K	35.97 M	
BVTSubGroup	787.29 M	1.7 M	1.5 M	4.87 %	99.90 %	73.04 K	35.97 M	

Drill-Downs


You can view the information displayed on this top-level report for a specific device, group/sub-group by using the following drill-down links:

- **Device:** Shows the Peak Rate vs. Average Rate graph, Network Efficiency graph, Class Utilization Device Table, and the Top 10 Children Pie Chart and Table for the top-bandwidth consuming child traffic classes for the selected class on this PacketShaper appliance.

- **Group:** Shows the Peak Rate vs. Average Rate graph, Network Efficiency graph, Class Utilization Device Table, and the Top 10 Children Pie Chart and Table for the top-bandwidth consuming child traffic classes for all PacketShaper appliances in the group or sub-group for the selected traffic class.

Top 10 Children Pie Chart and Table

Details the top-bandwidth consuming child classes reported on an individual device for the selected traffic class during the reporting period. Because the pie chart and table show classes for an individual device, you may notice the same traffic class shown multiple times. For example, if you run the report for traffic class /Inbound in a group that contains three PacketShaper appliances and HTTP is the top traffic class on all three, HTTP will show up as three separate pie slices (one for each PacketShaper) and it will be listed three times on the table. To identify which PacketShaper corresponds to a specific entry, hover over the Child Traffic Class name in the table.



Drill-Downs

You can view the information displayed on this top-level report for a specific device and class by using the following drill-down links:

- **Child Traffic Class:** Displays a breakdown of the average and peak rates, network efficiency, and utilization details for the specific child class on a specific PacketShaper appliance.
- Displays the list of top listeners for the selected child class on a specific PacketShaper appliance.
- Displays the list of top talkers for the selected child class on a specific PacketShaper appliance.

VoIP Statistics Report

The Voice over Internet Protocol (VoIP) Statistics report displays information that helps you determine the quality of the user experience for the RTP- and RTCP-based VoIP applications that are running on your network. The quality of the user experience for VoIP traffic can be calculated based on three statistics:

- **Loss**—the percentage of lost packets
- **Latency**—the time required for packets to travel from one PacketShaper to another, as measured by the formula round-trip time / 2. Latency is calculated on the Inbound interface only and can only be calculated between PacketShapers (one to intercept the inbound traffic at each end of a call). Latency information can only be collected from PacketShapers running PacketWise version 7.3.1 or later. If you are running an earlier version of PacketWise on your PacketShapers, you will not be able to see latency statistics for VoIP traffic running through those PacketShapers.
- **Jitter**—the variation in the delay of received packets in a flow


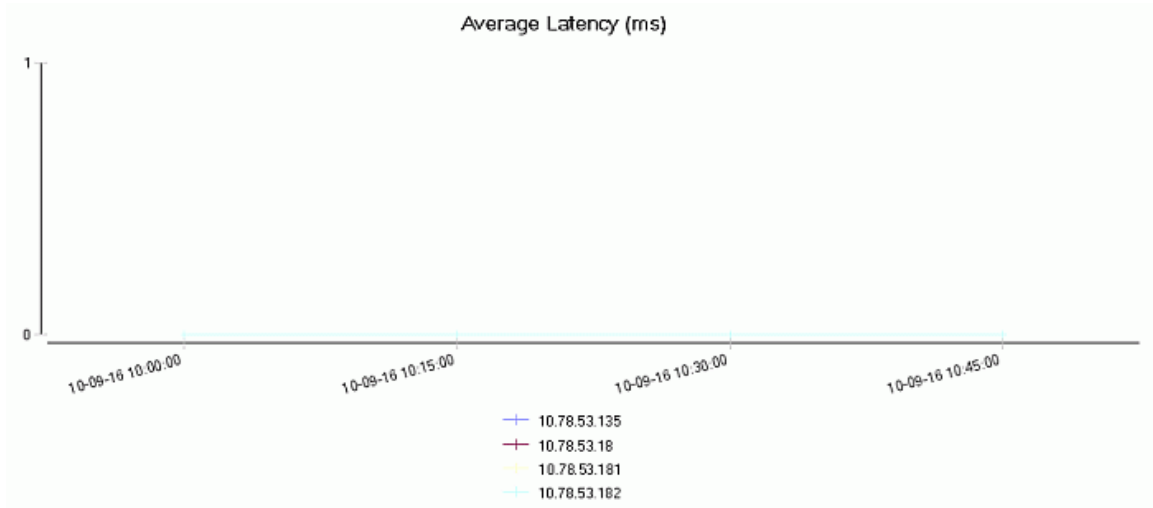
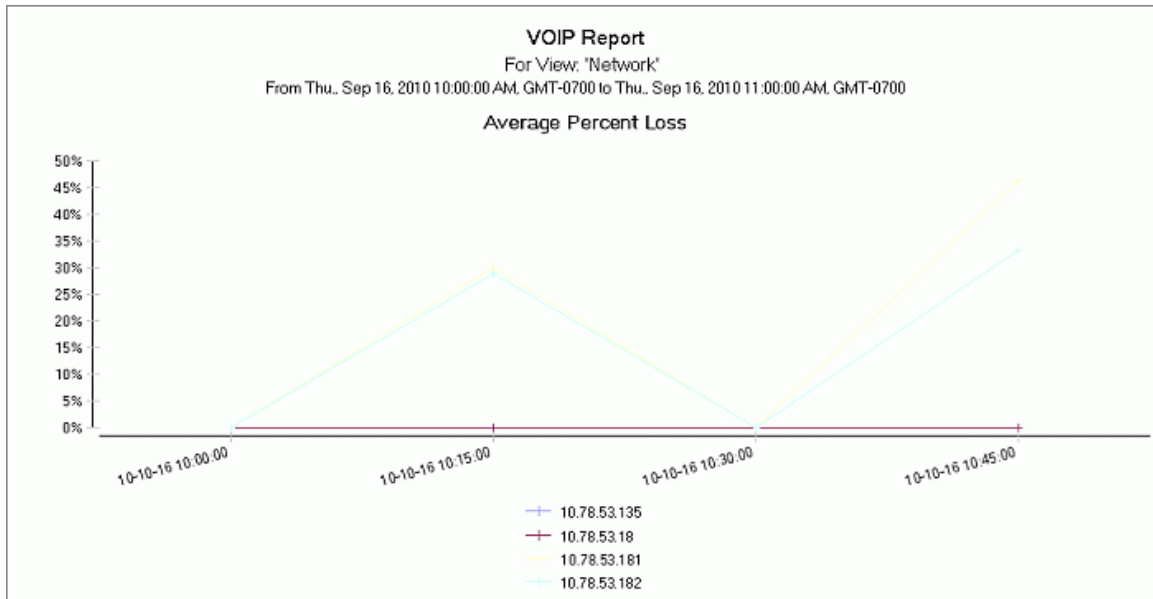
Additionally, DC uses the statistics it collects to calculate a mean opinion score (MOS) for a specific traffic class. The MOS value ranges from 1 (worst) to 5 (best) and provides a relative measure of VoIP quality.

IntelligenceCenter provides a VoIP application that includes all RTP-I, RTP-B, RTCP-I, and RTCP-B traffic; this application is selected by default when you configure the report. However, each VoIP application, device, or service uses a specific codec (coder/decoder) for the conversion between analog and digital signals. When you run the VoIP report, you must select the codec used by the VoIP application(s) that generated the traffic. The VoIP application is pre-selected and you must select the codec used by the VoIP application(s) that generated the traffic. Therefore, before you run the VoIP Statistics report, you should [modify the VoIP application definition](#) (**Configure > Application**) so that it contains VoIP applications that use the same codec. To determine which codec to use, refer to the documentation provided by the VoIP application, device, or service provider. Selecting the correct codec ensures the most accurate reporting and calculation of VoIP metrics such as percent loss, latency, jitter, and MOS.

This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.

Top-Level Views

VoIP Statistics Table and Graphs — Shows the IP address of each PacketShaper that collected VoIP traffic for the selected VoIP application and network group, sub-group, device or view. For each device, the table shows the percent loss and average latency (in milliseconds) and jitter (in milliseconds) for the VoIP calls that ran during the selected reporting period. These values are also shown on the three graphs below the table. To drill down to detailed information about the VoIP traffic classes reported by a specific PacketShaper, click on its IP address.



Device	Percent Loss	Average Latency	Average Jitter
10.78.53.135	0.00 %	0.00	0.00
10.78.53.18	0.00 %	0.00	0.00
10.78.53.181	0.00 %	0.00	219.34
10.78.53.182	0.00 %	0.00	199.87



Page: 1


Created on: Sep 16, 2010 11:36 AM

Drill Down: VoIP Statistics Per Device

This level of the report shows each traffic class that contributed to the VoIP traffic on the selected PacketShaper and details the loss, latency, and jitter values for each class. Additionally, the report shows graphs that detail the aggregated loss, latency, and jitter values for all classes reported on the device.


To drill down to the next level of the report, click a traffic class.






Drill Down: VoIP Statistics Per Traffic Class

When you drill down on a specific traffic class, you see graphs that show loss, latency, and jitter statistics for the selected class over the reporting period. In addition, a fourth graph shows the MOS value, which provides a relative measure of VoIP quality in the range of 1 (worst) to 5 (best).





Host Reports

Host Pairs Activity Report

The Host Pairs Activity report allows you to view information about all flows to and from a specific host running through a given PacketShaper appliance. Typically you would run this report to investigate suspicious network activity on a specific host. For example, you may notice on the [Top Host Pairs](#) or [Top Talkers](#) report that there is an abnormal amount of traffic being generated by a specific host. Or you may have an employee that you suspect is using network resources inappropriately. In this case, you could use the Host Pairs Activity report to view detailed information about the flows, traffic classes, and applications that the host is using.




When you configure the report, you specify the IP address of the host you are investigating.

Top-Level View



The top-level of this report shows a list of all flows to and from the specified host. This level of the report is generated based on Packeteer-2 FDR data. If you are not collecting Packeteer-2 data, you will not be able to generate this report.

For each host pair, this table shows the following information:

Statistic	Description
Addr 1 / Addr 2	The IP addresses of the host pair. The Host IP Address you specified when you configured the report to run will always be either a talker or a listener in each host pair.
Bytes	The total number of bytes used by the host pair
Flows	The total number of flows generated by the host pair
Packets	The total number of packets in the flows generated by the host pair
Efficiency	The percentage of bytes that did not require retransmission
Avg Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Avg Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the host pair detected during the report period

Host Pairs Activities Report							
For View: 'Network'							
From Mon, Oct 11, 2010 02:00:00 PM, GMT-0700 to Mon, Oct 11, 2010 03:00:00 PM, GMT-0700							
Addr 1 / Addr 2		Bytes	Flows	Packets	Efficiency	Avg Total Delay	Transactions
24.6.51.177 <-> 171.64.64.37		859.25 M	6	757.69 K	100.00 %	0.00	0
 To Host Pair Activity By Application		 To Host Pair Activity By Class					
Page: 1		Created on: Oct 11, 2010 3:06 PM					

Three drill-down reports are available:

- Clicking an IP address pair displays [flow details for the host pair](#).
- Clicking the  icon lists all the [host pair's applications](#). To drill-down to this level of the report you must also be collecting ME data.
- Clicking the  lists all the [host pair's traffic classes](#). To drill-down to this level of the report you must also be collecting ME data.

Top DSCP Report

The Top DSCP report enables you to identify the top Differentiated Services Code Point (DSCP) values that the traffic on your network is marked with.


When you configure the report, you specify how to determine the top DSCP values — by number of bytes or packets — as well as the number of top DSCP values to identify. This report is generated based on Packeteer-2 FDR data. If you are not collecting Packeteer-2 data, you will not be able to generate this report.

Top-Level Views

Pie Chart — Shows the bandwidth consumption of each top DSCP values relative to the others. Each pie slice represents the percentage of bandwidth each DSCP value comprised during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates each top DSCP value with a pie slice color.

Utilization Table — For each of the top DSCP values, this table shows the following information:

Statistic	Description
DSCPs	The DSCP value (based on the value of the ServiceType field in the Packeteer-2 packet)
Percent	The percentage of total bandwidth consumed by flows with the specified DSCP value relative to flows with the other top DSCP values during the reporting period
Bytes	The total number of bytes consumed by flows with this DSCP value
Flows	The total number of flows with this DSCP value
Packets	The total number of packets in the flows with the specified DSCP value
Peers	The total number of talkers with the specified DSCP value
Efficiency	The percentage of bytes that did not require retransmission
Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data with the specified DSCP value detected during the report period



Clicking on a DSCP value allows you to drill down to details about the hosts that generated the traffic with the selected value.

Drill Down: DSCP Host Details

By clicking on a DSCP value in the DSCP report, you can drill down to details about the hosts that generated the traffic with the selected value.

Top DSCP Table — For the specific DSCP value you selected at the top level of the report, this table displays a list of the hosts that sent traffic marked with the DSCP value during the reporting period. For each host entry, the table displays statistics for traffic that the host sent that was marked with the selected DSCP value.

Top DSCP: 0								
For View: 'Network'								
From Wed, Sep 01, 2010 12:00:00 AM, GMT-0700 to Fri, Sep 17, 2010 12:00:00 AM, GMT-0700								
Applications	Bytes	Flows	Packets	Efficiency	Avg Total Delay	Avg Server Delay	Transactions	
SSL	5.39 G	6.3 K	17.62 M	54.79	4.76 K	3.55 K	56.05 K	
QIFS	4.5 G	4.6 K	6.81 M	51.54	0.00	0.00	0	
SMTP	2.85 G	4.58 K	13.76 M	58.82	0.00	0.00	0	
HTTP	2.46 G	8.89 K	6.17 M	76.87	48.46	12.14	6.85 K	
ICMP	291.58 M	2.09 K	4.86 M	100.00	0.00	0.00	0	
DNS	13.12 M	6.06 K	138.77 K	100.00	0.00	0.00	0	
SSL-No-Cert	5.77 K	4	43	100.00	274.67	164.67	3	
lockd	1.22 K	16	16	100.00	0.00	0.00	0	
Telnet	652	2	16	100.00	181.67	0.33	3	
SNMP	630	9	9	100.00	0.00	0.00	0	
Echo	168	1	4	100.00	0.00	0.00	0	

Sep 17, 2010 1:31 PM

Clicking on an application allows you to drill down to flow details for the selected application.

Drill Down: DSCP Flow Details

In the DSCP host report, you can drill down to flow details for a host by clicking on the peer IP address or host name. Because the drill down shows actual flows, you will only see data for periods during which a flow with the selected DSCP value was occurring on the host. The first time you drill down on a host, you must specify the size of the interval IC uses to display the flow data: 5 minutes, 15 minutes, 30 minutes (default), 1 hour, or 2 hours. IC will then display the flow data for the selected interval, which you can scroll through using the **Next** or **Previous** links to view all flows within the reporting period. Note that flow data is stored as raw data and it can therefore only be displayed if it is still present in your database, which means that your reporting period must be within the length of time that your database is configured to store raw data. By default, DataCollector is configured to store raw data for 48 hours, but this value is [configurable](#).

DSCP-Host Table — For the host you selected on the DSCP host report, this table displays information about each flow that the host initiated with the specified DSCP value during the reporting period. Note that you will only be able to see flow data if your reporting period is within the window of time that raw data is stored in your database (48 hours by default). Additionally, because flows could have occurred at any time during the reporting period, each screen displays intervals of data and you will only see flow data if an actual flow occurred during the interval. To see the next interval, click **Next**.

Devices for DSCP: 0							
For View: 'Network'							
From Wed., Sep 15, 2010 12:00:00 PM, GMT-0700 to Wed., Sep 15, 2010 01:59:59 PM, GMT-0700							
Previous							
Next							
Device Name	Traffic Class	Bytes	Packets	Efficiency	Total Delay	Server Delay	Transactions
10.78.53.18-S	/Inbound/Localhost/SSL	34.25 K	189	100.00 %	2.81 K	1.56 K	26
10.78.53.18-S	/Inbound/SSL	16.46 M	66.98 K	77.37 %	0	0	0
10.78.53.18-S	/Outbound/Localhost/SSL	54.7 K	176	100.00 %	0	0	0
10.78.53.18-S	/Outbound/SSL	25.31 M	70.42 K	49.61 %	1.51 M	1.17 M	246
Previous							
Next							
Page: 1	Created on: Sep 17, 2010 1:33 PM						

Top Host Pairs Report

The Top Host Pairs report enables you to identify the top hosts that have exchanged the greatest amount of data on your network.


When you configure the report, you specify how to determine the top host pairs — by number of bytes or packets — as well as the number of top host pairs to identify. This report is generated based on Packeteer-2 FDR data. If you are not collecting Packeteer-2 data, you will not be able to generate this report.

Top-Level Views

Pie Chart — Shows the bandwidth consumption of each top host pair relative to the others. Each pie slice represents the percentage of bandwidth the host pair consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates the IP addresses of each top host pair with a pie slice color.

Utilization Table — For each of the top host pairs, this table shows the following information:

Statistic	Description
Addr 1 / Addr 2	The IP addresses of the host pair
Percent	The percentage of total bandwidth consumed by the host pair relative to the other top host pairs during the reporting period
Bytes	The total number of bytes used by the host pair
Flows	The total number of flows generated by the host pair
Packets	The total number of packets in the flows generated by the host pair
Efficiency	The percentage of bytes that did not require retransmission
Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the host pair detected during the report period



Clicking the IP address pair allows you to [drill down](#) to the actual flows for the host pair.

Top Listeners Report


The Top Listeners report enables you to identify the top hosts that receive the most traffic on your network.


When you configure the report, you specify how to determine the top listeners — by number of bytes or packets — as well as the number of top listeners to identify. This report is generated based on Packeteer-2 FDR data. If you are not collecting Packeteer-2 data, you will not be able to generate this report.

Top-Level Views


Pie Chart — Shows the bandwidth consumption of each top listener relative to the others. Each pie slice represents the percentage of bandwidth the listener consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates each top listener's IP address with a pie slice color.

Utilization Table — For each of the top listeners, this table shows the following information:

Statistic	Description
Listeners	The IP address or host name of the listener
	Drill down to Top Applications for the listener
Percent	The percentage of total bandwidth consumed by the host relative to the other top listeners during the reporting period
Bytes	The total number of bytes used by the host
Flows	The total number of flows generated by the host
Packets	The total number of packets in the flows generated by the host
Peers	Number of peers the listener has paired with
Efficiency	The percentage of bytes that did not require retransmission
Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the host detected during the report period



This report offers two drill-down reports:

- [Top Listener Peers](#) — Clicking on an IP address allows you to drill down to details about the specific host pairs for the listener.
- [Top Applications for Listener](#) — Clicking the  next to the IP address allows you to drill down to a list of top applications for a particular top listener.

Top Services Report

Knowing the identity of traffic running over your network is a big first step in managing and controlling the performance of network applications. The Top Services report enables you to identify the top services that are running on your network. When you configure the report, you specify how to determine the top services — by number of bytes or packets — as well as the number of top services to identify. This report is generated based on Packeteer-2 FDR data. You must collect Packeteer-2 data to generate this report.


Top-Level Views

Pie Chart — Shows the bandwidth consumption of each top service relative to the others. Each pie slice represents the percentage of bandwidth the service consumed during the reporting period. Hover over a slice to see specific percentages. The legend associates a service name with a pie slice color.

Service Utilization Table — For each of the top services, this table shows utilization and response time statistics. Clicking on a service name allows you to drill down to details about the specific devices and groups responsible for generating the traffic flows for the selected service.

Statistic	Description
Service Name*	The name of the service (based on the value of the ServiceType field in the Packeteer-2 packet)
Percent*	The percentage of total bandwidth consumed by the service relative to the other top services during the reporting period
Bytes*	The total number of bytes consumed by flows for this service
Flows*	The total number of flows generated by the service
Packets*	The total number of packets in the flows generated by the service
Peers	The total number of service talkers
Efficiency	The percentage of bytes that did not require retransmission
Avg Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Avg Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the service detected during the report period

* Sortable column



Drill Down: Service Host List

By clicking on a service name in the Top Services report, you can drill down to details about the specific devices and groups responsible for generating the traffic flows for the selected service.

Traffic utilization table for a specific service — For the specific service you selected at the top level of the report, this table displays a list of the hosts that generated flows for that service during the reporting period. For each host entry, the table displays the same statistics per-host for the selected host that were shown at the top level of the report for the selected service generated by each host.

Top Service: HTTP							
For View: 'Network'							
From Thu., Sep 16, 2010 12:00:00 AM, GMT-0700 to Fri., Sep 17, 2010 12:00:00 AM, GMT-0700							
Hosts	Bytes	Flows	Packets	Efficiency	Avg Total Delay	Avg Server Delay	Transactions
10.1.61.5	992.37 M	434.65 K	2.33 M	72.87	0.00	0.00	0
10.1.61.13	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.4	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.17	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.1	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.28	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.9	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.3	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.10	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.24	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.26	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.30	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.14	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.61.7	951.13 M	412.29 K	2.23 M	73.08	0.00	0.00	0
10.1.41.12	935.33 M	402.32 K	2.16 M	71.74	0.00	0.00	0
10.1.41.28	935.33 M	402.32 K	2.16 M	71.74	0.00	0.00	0
10.1.41.5	935.33 M	402.32 K	2.16 M	71.74	0.00	0.00	0
10.1.41.9	935.33 M	402.32 K	2.16 M	71.74	0.00	0.00	0
10.1.41.4	935.33 M	402.32 K	2.16 M	71.74	0.00	0.00	0
10.1.41.26	935.33 M	402.32 K	2.16 M	71.74	0.00	0.00	0
10.1.41.29	935.33 M	402.32 K	2.16 M	71.74	0.00	0.00	0
10.1.41.16	935.33 M	402.32 K	2.16 M	71.74	0.00	0.00	0
10.1.41.1	935.33 M	402.32 K	2.16 M	71.74	0.00	0.00	0

Clicking on a specific host entry allows you to drill down to the Average Rate - Peak Rate Comparison for the device or group.

Drill Down: Flow Details

To drill down to flow details for a particular host in the Top Service report's traffic utilization table, you can click on a peer IP address or host name. Because the second-level drill down shows actual flows, you will only see data for periods during which a flow with the selected service value was occurring on the host. The first time you drill down on a host, you must specify the size of the interval IC uses to display the flow data: 5 minutes, 15 minutes, 30 minutes (default), 1 hour, or 2 hours. IC will then display the flow data for the selected interval, which you can scroll through using the **Next** or **Previous** links to view all flows within the reporting period. Note that flow data is stored as raw data and it can therefore only be displayed if it is still present in your database, which means that your reporting period must be within the length of time that your database is configured to store raw data. By default, DataCollector is configured to store raw data for 48 hours, but this value is [configurable](#).

Service flows — For the host you selected on the traffic utilization table, this table displays information about each flow that the host initiated with the specified service type during the reporting period. Note that you will only be able to see flow data if your reporting period is within the window of time that raw data is stored in your database (48 hours by default). Additionally, because flows could have occurred at any time during the reporting period, each screen displays intervals of data and you will only see flow data if an actual flow occurred during the interval. To see the next interval, click **Next**. For each flow, the table shows the start time, end time, and

duration and throughput of each flow, the IP address and port of the source and destination hosts, the protocol, service, the traffic class into which the flow was classified, policy, and efficiency percentage for each flow.

Service: HTTP with Host: 10.1.61.5										
For View: 'Network'										
From Thu., Sep 16, 2010 12:00:00 AM, GMT-0700 to Thu., Sep 16, 2010 01:59:59 AM, GMT-0700 ^										
Next										
Start Time	End Time	Duration	Throughput	Src Ip	Src Port	Dest Ip	Dest Port	Service	Traffic Class	Efficiency
2010-09-16 01:36:37-0700	2010-09-16 01:36:51-0700	14	1817	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:36:37-0700	2010-09-16 01:36:51-0700	14	954	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:36:57-0700	2010-09-16 01:37:12-0700	15	1696	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:36:58-0700	2010-09-16 01:37:12-0700	14	954	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:37:15-0700	2010-09-16 01:37:30-0700	15	1696	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:37:16-0700	2010-09-16 01:37:30-0700	14	954	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:37:35-0700	2010-09-16 01:37:50-0700	15	1696	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:37:35-0700	2010-09-16 01:37:50-0700	15	890	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:37:54-0700	2010-09-16 01:38:08-0700	14	1817	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:37:54-0700	2010-09-16 01:38:08-0700	14	954	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:38:13-0700	2010-09-16 01:38:27-0700	14	1817	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:38:13-0700	2010-09-16 01:38:27-0700	14	954	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:38:31-0700	2010-09-16 01:38:46-0700	15	1696	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:38:31-0700	2010-09-16 01:38:46-0700	15	890	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:38:51-0700	2010-09-16 01:39:05-0700	14	1817	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:38:51-0700	2010-09-16 01:39:05-0700	14	954	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:39:08-0700	2010-09-16 01:39:24-0700	16	1590	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:39:09-0700	2010-09-16 01:39:24-0700	15	890	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:39:30-0700	2010-09-16 01:39:43-0700	13	1956	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:39:30-0700	2010-09-16 01:39:43-0700	13	1027	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:39:47-0700	2010-09-16 01:40:01-0700	14	1817	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:39:47-0700	2010-09-16 01:40:01-0700	14	954	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %
2010-09-16 01:40:08-0700	2010-09-16 01:40:22-0700	14	1817	10.1.61.5	43956	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	52.00 %
2010-09-16 01:40:08-0700	2010-09-16 01:40:22-0700	14	954	10.1.61.5	43957	10.2.61.4	80	HTTP	/Outbound/Internet/HTTP	100.00 %

Top Talkers Report


The Top Talkers report enables you to identify the top hosts that initiate the most traffic on your network.


When you configure the report, you specify how to determine the top talkers — by number of bytes or packets — as well as the number of top talkers to identify. This report is generated based on Packeteer-2 FDR data. If you are not collecting Packeteer-2 data, you will not be able to generate this report.

Top-Level Views


Top Pie Chart — Shows the bandwidth consumption of each top talker relative to the others. Each pie slice represents the percentage of bandwidth the talker consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates each top talker's IP address with a pie slice color.

Utilization Table — For each of the top talkers, this table shows the following information:

Statistic	Description
Talkers	The IP address or host name of the talker.
	Drill down to Top Applications for the talker; you must be collecting ME data to drill down to this report level.
Percent	The percentage of total bandwidth consumed by the host relative to the other top talkers during the reporting period.
Bytes	The total number of bytes used by the host.
Flows	The total number of flows generated by the host.
Packets	The total number of packets in the flows generated by the host.
Peers	Number of peers the talker has paired with.
Efficiency	The percentage of bytes that did not require retransmission.
Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay.
Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the host detected during the report period.



This report offers two drill-down reports:

- [Top Talker Peers](#) — Clicking on an IP address allows you to drill down to details about the specific host pairs for the talker.
- [Top Applications for Talker](#) — Clicking the  next to the IP address allows you to drill down to a list of top applications for a particular top talker. You must be collecting ME data to drill down to this report level.

Top VLAN Report

The Top VLAN report enables you to identify the top VLAN IDs on your network.


When you configure the report, you specify how to determine the top VLAN IDs — by number of bytes or packets — as well as the number of top VLAN IDs to identify. This report is generated based on Packeteer-2 FDR data. If you are not collecting Packeteer-2 data, you will not be able to generate this report.

Top-Level Views

Pie Chart — Shows the bandwidth consumption of each top VLAN ID relative to the others. Each pie slice represents the percentage of bandwidth the VLAN ID comprised during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates each top VLAN ID with a pie slice color.

Utilization Table — For each of the top VLAN IDs, this table shows the following information:

Statistic	Description
VLANs	The value associated with the VLAN
Percent	The percentage of total bandwidth consumed by flows with the specified VLAN ID relative to flows with the other top VLAN IDs during the reporting period
Bytes	The total number of bytes consumed by flows with this VLAN ID
Flows	The total number of flows with this VLAN ID
Packets	The total number of packets in the flows with the specified VLAN ID
Peers	The total number of talkers with the specified VLAN ID
Efficiency	The percentage of bytes that did not require retransmission
Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data with the specified VLAN ID detected during the report period



Click on a VLAN ID to drill down to details about the hosts that generated the traffic with the selected value.

Drill Down: VLAN Host Details

In the Top VLAN report, you can click on a VLAN ID to drill down to details about the hosts that generated the traffic with the selected value.

Top Hosts for a specific VLAN ID — For the specific VLAN ID you selected at the top level of the report, this table displays a list of the hosts that sent traffic marked with the VLAN ID during the reporting period. For each host entry, the table displays the total number of packets, bytes, and flows that the host sent that was marked with the selected VLAN ID.

Top VLAN: 0							
For View: 'Network'							
From Wed., Sep 01, 2010 12:00:00 AM, GMT-0700 to Fri., Sep 17, 2010 12:00:00 AM, GMT-0700							
Applications	Bytes	Flows	Packets	Efficiency	Avg Total Delay	Avg Server Delay	Transactions
SSL	5.37 G	6.28 K	17.55 M	54.73	4.76 K	3.56 K	55.9 K
CIFS	4.49 G	4.6 K	6.79 M	51.51	0.00	0.00	0
SMTP	2.84 G	4.57 K	13.71 M	58.77	0.00	0.00	0
HTTP	2.46 G	8.87 K	6.16 M	76.89	48.05	12.14	6.83 K
ICMP	291.58 M	2.09 K	4.86 M	100.00	0.00	0.00	0
DNS	13.07 M	6.05 K	138.27 K	100.00	0.00	0.00	0
SSL-No-Cert	5.77 K	4	43	100.00	274.67	164.67	3
lockd	1.22 K	16	16	100.00	0.00	0.00	0
Telnet	652	2	16	100.00	181.67	0.33	3
SNMP	630	9	9	100.00	0.00	0.00	0
Echo	168	1	4	100.00	0.00	0.00	0

Sep 17, 2010 2:23 PM

Click on a peer IP address or host name to drill down specific flow details for the specific flows.

Drill Down: VLAN Flow Details

In the VLAN Host report, you can click on a peer IP address or host name to drill down specific flow details for the specific flows. Because the drill down shows actual flows, you will only see data for periods during which a flow with the selected VLAN value was occurring on the host. The first time you drill down on a host, you must specify the size of the interval IC uses to display the flow data: 5 minutes, 15 minutes, 30 minutes (default), 1 hour, or 2 hours. IC will then display the flow data for the selected interval, which you can scroll through using the **Next** or **Previous** links to view all flows within the reporting period. Note that flow data is stored as raw data and it can therefore only be displayed if it is still present in your database, which means that your reporting period must be within the length of time that your database is configured to store raw data. By default, DataCollector is configured to store raw data for 48 hours, but this value is [configurable](#).

Devices for VLAN: 0									
For View: 'Network'									
From Wed., Sep 15, 2010 02:00:00 PM, GMT-0700 to Wed., Sep 15, 2010 03:59:59 PM, GMT-0700									
Previous									
Next									
Device Name	Traffic Class	Interface In	Interface Out	Bytes	Packets	Efficiency	Total Delay	Server Delay	Transactions
10.78.53.10-S	/Inbound/Internet/HTTP	2	1	51.53 M	140.89 K	86.66 %	0	0	0
10.78.53.18-S	/Inbound	4	3	429.75 K	2.29 K	100.00 %	1.85 K	465	43
10.78.53.18-S	/Localhost/HTTP								
10.78.53.18-S	/Outbound	1	2	56.17 M	122.71 K	68.82 %	0	0	0
10.78.53.18-S	/Internet/HTTP								
10.78.53.18-S	/Outbound	3	4	117.3 K	1.63 K	100.00 %	0	0	0
10.78.53.18-S	/Localhost/HTTP								
Previous									
Next									
Page: 1									
Created on: Sep 17, 2010 2:24 PM									

Site Reports

Site Response Time Report

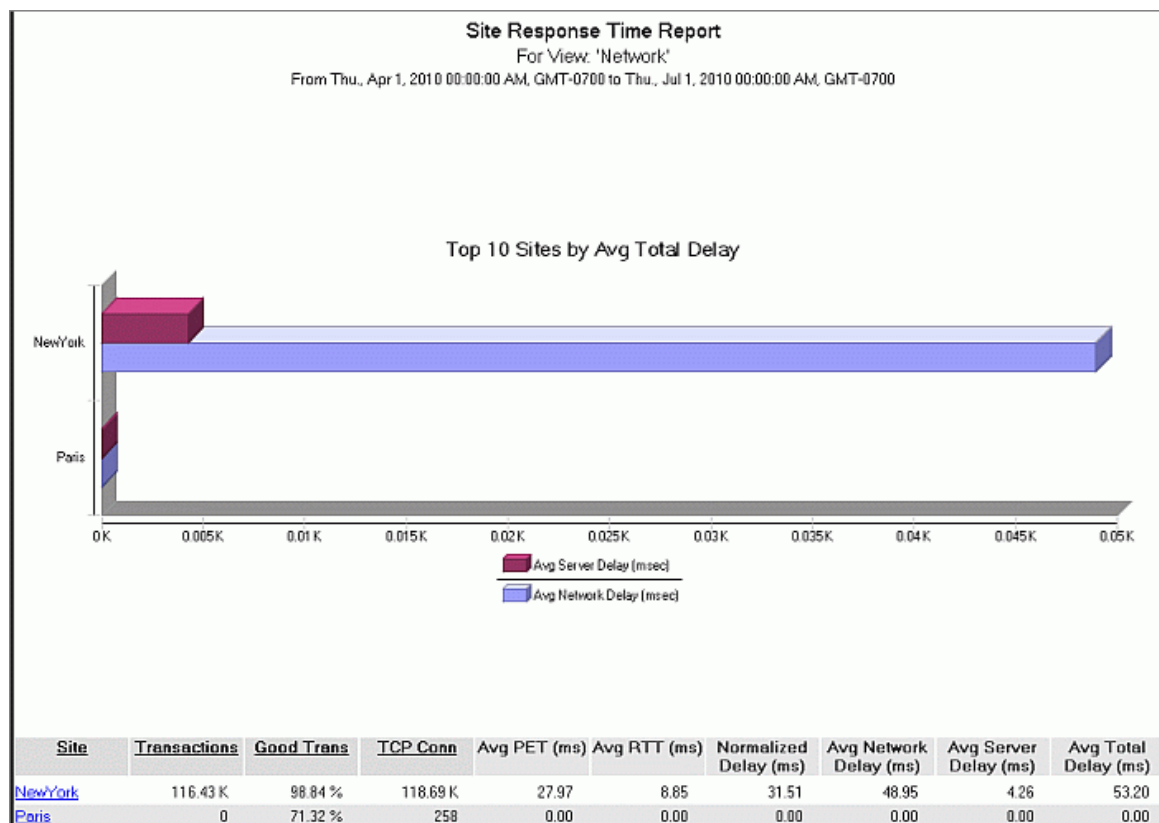
Response time measurement (RTM) provides information about the amount of time connection-based TCP traffic spends traveling between a client and a server and the time used by the server itself. This allows you to investigate response times and identify the source of network delays. The Site RTM report contrasts RTM statistics for the TCP traffic on each site.

When you configure the report, you specify the network group, sub-group, device, or view to which to restrict the report. For example, you could restrict the report so that it shows RTM statistics for a specific branch office or within a specific department.

This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.

Top-Level View

Top 10 Sites by Total Delay Graph — Compares the network delay and server delay of connection-based TCP traffic at the top 10 sites in the selected network group.



RTM Table — Shows response time statistics for each site in the selected network group. This table contains the following information:

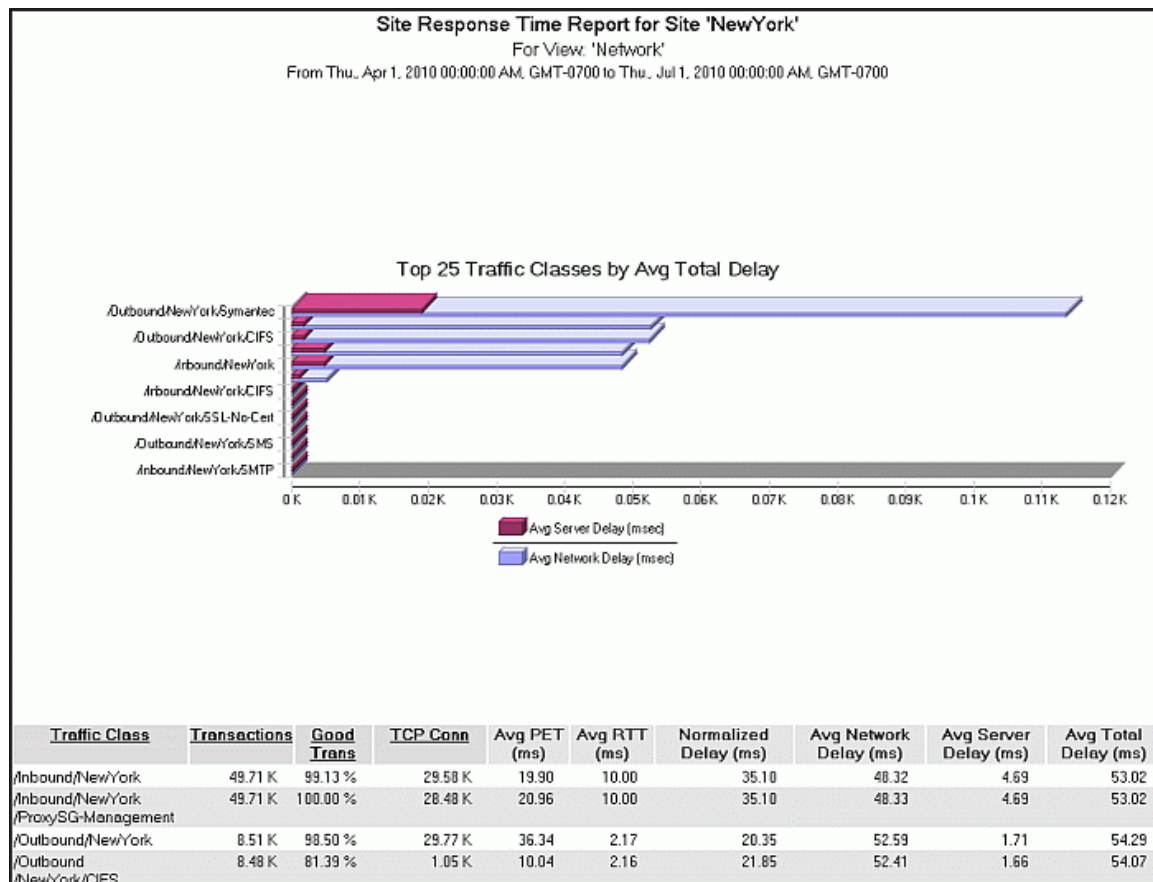
Statistic	Description
Transactions	The total number of transactions reported for the application
Good Trans	The percentage of transactions that completed within the total delay threshold
TCP Conn	The number of TCP connections established at the site and used by the transactions that were counted in the summary
Avg PET (ms)	The average Packet Exchange Time (PET) for transactions at the site. The PET represents the interval between a data packet leaving the PacketShaper and its acknowledgment (ACK) arriving.
Avg RTT (ms)	The average Round Trip Time (RTT) for packets at the site. The RTT represents the average number of milliseconds spent in transit when a client and server exchange the SYN (synchronize sequence numbers flag) and its corresponding ACK (acknowledge flag). A transaction involving a large amount of data requires the data to be divided into multiple packets. Whereas a transaction's network delay reflects the total transit time for all required packets, the RTT reflects the time for a single small packet to make its way from client to server and another packet to make the return trip. Use the RTT to determine if a large network delay is due to large transactions or a slow network. If the RTT is much smaller than the network delay, the transactions were large. If the two averages are close, a sluggish network caused the longer network delays.
Normalized Delay (ms)	The transaction delay in the network, normalized by transaction size. It shows how long it takes to send 1KB of data. This statistic allows an accurate comparison of response-time data for different applications or servers. Without normalizing the delay, response times vary depending on the size of the transaction. This statistic eliminates size as a factor of network delay
Avg Network Delay (ms)	The average response times in milliseconds of the TCP traffic in the site over the reporting period. This statistic represents only the portion of the transaction time that is attributable to the network, enabling you to analyze network delay.
Avg Server Delay (ms)	The average response times in milliseconds of the TCP traffic in the site over the reporting period. This graph shows only the portion of the transaction time that is attributable to the server, enabling you to analyze server delay.
Avg Total Delay (ms)	Average number of milliseconds to complete transactions; includes network delay and server delay

Clicking on a site name allows you to drill down to see the RTM data for traffic classes at the selected site.

Drill Down: Traffic Class RTM for Site

By clicking on a site name in the Site RTM report, you can drill down to see the per-class RTM data at the selected site. Note that if the site spans multiple PacketShaper appliances, the data from all of the appliances will be aggregated for the site.

Traffic Class RTM for a site — For the selected site, details the RTM statistics for each traffic class that reported RTM data during the reporting period. The table shows the same RTM statistics that were displayed on the top-level view for the site as a whole.



Top Applications by Site Report


To get an understanding of what applications are using the most bandwidth at a particular site during a specific time frame, you can run the Top Applications by Site report.

When you configure the report, you select the network group and site, specify how to determine the top applications — by number of bytes or packets — as well as the number of top applications to identify. This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.


Top-Level Views

Pie Chart — Shows the bandwidth consumption of each top application relative to the others at the selected site. Each pie slice represents the percentage of bandwidth each of the applications consumed during the reporting period. Hover over a slice to see the associated application name and summary details. The legend to the right of the pie chart associates an application name with a pie slice color. Note that if the site spans multiple PacketShaper appliances, the data from all of the appliances will be aggregated for the site.

Application Utilization Table — For each of the top applications, this table shows the ranking of the application in bandwidth consumption, the name of the application, the percentage of total bandwidth consumed by the application, the total number of packets and bytes used by the application, the efficiency level percentage, the average and peak rates for the application, and number of guaranteed rate failures (if this option was selected when running the report) for the reporting period. Clicking an application name allows you to [drill down](#) to details about the specific devices and groups responsible for the application traffic.



Average Rate Stack Chart — Shows throughput in the network for the top applications.



Top Host Pairs by Site Report

The Top Host Pairs report enables you to identify the pairs of hosts that have exchanged the greatest amount of data at a particular site.


When you configure the report, you select the network group and site, specify how to determine the top host pairs — by number of bytes or packets — as well as the number of top host pairs to identify. This report is generated based on Packeteer-2 FDR data. If you are not collecting Packeteer-2 data, you will not be able to generate this report.

Top-Level Views

Pie Chart — Shows the bandwidth consumption of each top host pair relative to the others at the selected site. Each pie slice represents the percentage of bandwidth the host pair consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates the IP addresses of each top host pair with a pie slice color. Note that if the site spans multiple PacketShaper appliances, the data from all of the appliances will be aggregated for the site.

Utilization Table — For each of the top host pairs, this table shows the following information:

Statistic	Description
Addr 1 / Addr 2	The IP addresses of the host pair
Percent	The percentage of total bandwidth consumed by the host pair relative to the other top host pairs during the reporting period
Bytes	The total number of bytes used by the host pair
Flows	The total number of flows generated by the host pair
Packets	The total number of packets in the flows generated by the host pair
Efficiency	The percentage of bytes that did not require retransmission
Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the host pair detected during the report period



Clicking on the IP address pair allows you to [drill down](#) to the actual flows for the host pair within the selected site.

Top Listeners by Site Report

The Top Listeners by Site report enables you to identify the hosts that receive the most traffic at a particular site.


When you configure the report, you select the network group and site, specify how to determine the top listeners — by number of bytes or packets — as well as the number of top listeners to identify. This report is generated based on Packeteer-2 FDR data. If you are not collecting Packeteer-2 data, you will not be able to generate this report.

Top-Level Views

Pie Chart — Shows the bandwidth consumption of each top listener relative to the others at the site. Each pie slice represents the percentage of bandwidth the listener consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates each top listener's IP address with a pie slice color. Note that if the site spans multiple PacketShaper appliances, the data from all of the appliances will be aggregated for the site.

Utilization Table — For each of the top listeners, this table shows the following information:

Statistic	Description
Listeners	The IP address or host name of the listener
Percent	The percentage of total bandwidth consumed by the host relative to the other top listeners during the reporting period
Bytes	The total number of bytes used by the host
Flows	The total number of flows generated by the host
Packets	The total number of packets in the flows generated by the host
Peers	Number of peers the listener has paired with
Efficiency	The percentage of bytes that did not require retransmission
Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the host detected during the report period



Clicking on an IP address allows you to [drill down](#) to details about the specific host pairs for the listener within the selected site.

Top Services by Site

Knowing the identity of traffic running over your network is a big first step in managing and controlling the performance of network applications. The Top Services by Site report enables you to identify the top services that are running at a particular site.

When you configure the report, you specify how to determine the top services — by number of bytes or packets — as well as the number of top services to identify. This report is generated based on Packeteer-2 FDR data. If you are not collecting Packeteer-2 data, you will not be able to generate this report.


Top-Level Views

Pie Chart — Shows the bandwidth consumption of each top service relative to the others. Each pie slice represents the percentage of bandwidth the service consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates a service name with a pie slice color.

Service Utilization Table — For each of the top services, this table shows utilization and response time statistics. Clicking on a service name allows you to drill down to details about the specific devices and groups responsible for generating the traffic flows for the selected service.

Statistic	Description
Services*	The name of the service (based on the value of the ServiceType field in the Packeteer-2 packet)
Percent*	The percentage of total bandwidth consumed by the service relative to the other top services during the reporting period
Bytes*	The total number of bytes consumed by flows for this service
Flows*	The total number of flows generated by the service
Packets*	The total number of packets in the flows generated by the service
Efficiency	The percentage of bytes that did not require retransmission
Avg Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Avg Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the service detected during the report period

* Sortable column



To drill down to a report detailing the host pairs that generated traffic for a particular service, click the service name.

Host Pairs
 For View: 'Network'
 For Site: SubnetClass1
 For Service: HTTP
 From Wed., Dec 08, 2010 10:00:00 AM, GMT-0800 to Wed., Dec 08, 2010 11:00:00 AM, GMT-0800

Addr1 / Addr2	Bytes	Packets	Flows	Efficiency	Avg Total Delay	Avg Server Delay	Transactions
10.9.68.144 <=> 10.125.32.23	2.95 M	6.38 K	84	1.00 %	0.00	0.00	0
10.125.4.10 <=> 10.125.32.23	654.55 K	4.52 K	114	0.99 %	503.63	367.31	300
10.9.68.145 <=> 10.125.32.23	108.99 K	1.13 K	187	1.00 %	0.00	0.00	0
10.125.32.11 <=> 10.125.32.10	476	4	2	1.00 %	0.00	0.00	0

Page: 1 Created on: Dec 8, 2010 11:43 AM

Top Sites Report

To compare the amount of traffic utilized on each of the sites on your network, use the Top Sites report. This report is a launching point for viewing other site-based reports; for each top site, you can drill down to see its top applications, traffic classes, talkers, listeners, and host pairs.


When you configure the report, you specify how to determine the top sites — by number of bytes or packets — as well as the number of top sites to identify.

This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report. In addition, some of the drill-down reports available from this report require that you are also collecting Packeteer-2 FDR data. For the data between report levels to correlate, you must be collecting ME and FDR from the same set of appliances.


Top-Level Views

Pie Chart — Shows the bandwidth consumption of each site relative to the other top sites. Each pie slice represents the percentage of bandwidth the site consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates a site name with a pie slice color. Note that if the site spans multiple PacketShaper appliances, the data from all of the appliances will be aggregated for the site.

Site Utilization Table — For each of the top sites, this table shows the ranking of the site in terms of bandwidth consumption, the name of the site, the percentage of total bandwidth consumed at the site, the total number of packets and bytes used at the site, the efficiency level percentage, the average and peak rates for the site, and the number of guaranteed rate failures during the reporting period.



Average Rate Stack Chart — Shows throughput in the network for the top sites.



Drill Downs



The following drill-down reports are also available by clicking the color-coded symbol next to the site name:

Symbol	Drill-Down Report
	Top Applications by Site —This drill-down report is based on ME data.
	Top Talkers by Site —To drill down to the top talkers for the site, you must also be collecting Packeteer-2 FDR data. For the data between levels of the report to correlate, you must be collecting ME and FDR from the same set of appliances.
	Top Listeners by Site —To drill down to the top listeners for the site, you must also be collecting Packeteer-2 FDR data. For the data between levels of the report to correlate, you must be collecting ME and FDR from the same set of appliances.
	Top Host Pairs by Site Report —To drill down to the top host pairs for the site, you must also be collecting Packeteer-2 FDR data. For the data between levels of the report to correlate, you must be collecting ME and FDR from the same set of appliances.
	Top Services by Site —To drill down to the top host pairs for the site, you must also be collecting Packeteer-2 FDR data. For the data between levels of the report to correlate, you must be collecting ME and FDR from the same set of appliances.

You can also drill down on a site name to see more details about an individual site as follows:

Site Utilization Comparison Report					
For View: 'Network'					
For Site: 'SubnetClass1'					
From Wed., Dec 08, 2010 10:00:00 AM, GMT-0800 to Wed., Dec 08, 2010 11:00:00 AM, GMT-0800					
Site	Bytes	Packets	Efficiency	Avg Rate (bps)	Peak Rate (bps)
SubnetClass1	4.81 M	15.53 K	99.86 %	313.50	981.57 K

Click the site name to drill down to view graphs detailing site utilization:



Top Talkers by Site Report

The Top Talkers by Site report enables you to identify the hosts that initiate the most traffic at a particular site.


When you configure the report, you select the network group and site, specify how to determine the top talkers — by number of bytes or packets — as well as the number of top talkers to identify. This report is generated based on Packeteer-2 FDR data. If you are not collecting Packeteer-2 data, you will not be able to generate this report.

Top-Level Views

Pie Chart — Shows the bandwidth consumption of each top talker relative to the others at the site. Each pie slice represents the percentage of bandwidth the talker consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates each top talker's IP address with a pie slice color. Note that if the site spans multiple PacketShaper appliances, the data from all of the appliances will be aggregated for the site.

Utilization Table — For each of the top talkers, this table shows the following information:

Statistic	Description
Talkers	The IP address or host name of the talker
Percent	The percentage of total bandwidth consumed by the host relative to the other top talkers during the reporting period
Bytes	The total number of bytes used by the host
Flows	The total number of flows generated by the host
Packets	The total number of packets in the flows generated by the host
Peers	Number of peers the talker has paired with
Efficiency	The percentage of bytes that did not require retransmission
Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the host detected during the report period



Clicking on an IP address allows you to [drill down](#) to details about the specific host pairs for the talker within the selected site.

Drill-Down Reports

Drill Down: Application Activity

In the Top Applications reports (such as [Top Applications](#), [Top Applications by Site](#), and [Top Applications for Talker/Listener](#)), you can click on an application name to drill down to details about the specific devices and groups responsible for the application traffic. This report is generated based on ME data. If you are not collecting ME data, you will not be able to generate this report.


Application Utilization Table for a specific application — For the specific application you selected at the top level of the Top Applications report, this table displays a list of the devices and/or groups that contributed to the application traffic during the reporting period. For each device or group entry, the table displays the total number of packets and bytes used by the device or group, the efficiency level percentage, and the average and peak rates for the application traffic generated by the device or group.

Application Activity - Devices and Groups						
For Application: 'ICMP'						
From Fri, Sep 17, 2010 12:00:00 PM, GMT-0700 to Fri, Sep 17, 2010 01:00:00 PM, GMT-0700						
IP Address	Device Name	Bytes	Packets	Efficiency	Avg Rate (bps)	Peak Rate (bps)
10.78.53.181	Lab-PS10K-2	595.77 M	6.26 M	100.00 %	13.39 K	1.34 M
10.78.53.182	Lab-PS10K	594.91 M	6.25 M	100.00 %	13.35 K	1.34 M
10.78.53.18	10.78.53.18-S	12.4 M	200.07 K	100.00 %	6.89 K	369.01 K
10.78.53.135	Sunnyvale	8.3 M	32.29 K	100.00 %	3.07 K	147.6 K
Group	Bytes	Packets	Efficiency	Avg Rate (bps)	Peak Rate (bps)	
Network	1.21 G	12.74 M	100.00 %	12.95 K	1.34 M	
10.78.53.18-S	12.4 M	200.07 K	100.00 %	6.89 K	369.01 K	

Click on a specific device name or IP address or group name to [drill down](#) to the Average Rate - Peak Rate Comparison for the device or group.

Drill Down: Application Rates Comparison

Average Rate - Peak Rate Comparison — Shows a line chart that details the average rate versus the peak rate for the specific application-device or group combination you selected at the previous level of the report.



Drill Down: Host Pairs Activity by Application

When drilling down from the Host Pairs Activity report, the "By Application" report lists all the host pair's applications as well as the following details for each application:

Statistic	Description
Bytes	The total number of bytes used by the host pair for the application
Flows	The total number of flows generated by the host pair for the application
Packets	The total number of packets in the flows generated by the host pair for the application
Efficiency	The percentage of bytes that did not require retransmission for the application
Avg Total Delay	For the application, average number of milliseconds to complete transactions; includes network delay and server delay
Avg Server Delay	For the application, average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	For the application, the total number of unique iterations of a transfer of data for the host pair detected during the report period

Host Pairs Activities Report							
(By Application)							
Host Pair: 24.6.51.177 <=> 171.64.64.37							
From Mon., Oct 11, 2010 02:00:00 PM, GMT-0700 to Mon., Oct 11, 2010 03:00:00 PM, GMT-0700							
Application	Bytes	Packets	Flows	Efficiency	Avg Total Delay	Avg Server Delay	Transactions
ICMP	7.63 M	190.07 K	3	100.00 %	0.00	0.00	0
Page: 1							
Created on: Oct 11, 2010 3:20 PM							

Drill Down: Host Pairs Activity by Traffic Class

When drilling down from the Host Pairs Activity report, the "By Traffic Class" report lists all the host pair's traffic classes as well as the following details for each class:

Statistic	Description
Bytes	The total number of bytes used by the host pair for the traffic class
Flows	The total number of flows generated by the host pair for the traffic class
Packets	The total number of packets in the flows generated by the host pair for the traffic class
Efficiency	The percentage of bytes that did not require retransmission for the traffic class
Avg Total Delay	For the traffic class, average number of milliseconds to complete transactions; includes network delay and server delay
Avg Server Delay	For the traffic class, average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	For the traffic class, the total number of unique iterations of a transfer of data for the host pair detected during the report period

Host Pairs Activities Report							
(By Traffic Class)							
Host Pair: 24.6.51.177 <=> 171.64.64.37							
From Mon., Oct 11, 2010 02:00:00 PM, GMT-0700 to Mon., Oct 11, 2010 03:00:00 PM, GMT-0700							
Traffic Class	Bytes	Packets	Flows	Efficiency	Avg Total Delay	Avg Server Delay	Transactions
/Inbound/LA/Default	851.42 M	567.62 K	3	100.00 %	0.00	0.00	0
/Outbound/Auburn/ICMP	7.83 M	190.07 K	3	100.00 %	0.00	0.00	0
Page: 1	Created on: Oct 11, 2010 3:22 PM						

Drill Down: Host-Pair Flow Details

To drill down to flow details for a specific pair of hosts, you can click on a peer IP address or host name in the peer table included in any [Top Talker/Listener](#) or [Host Pair](#) report. Because this drill down shows actual flows, you will only see data for periods during which flows were exchanged between the selected pair. The first time you drill down on a host, you must specify the size of the interval IC uses to display the flow data: 5 minutes, 15 minutes, 30 minutes (default), 1 hour, or 2 hours. IC will then display the flow data for the selected interval, which you can scroll through using the **Next** or **Previous** links to view all flows within the reporting period. Note that flow data is stored as raw data and it can therefore only be displayed if it is still present in your database, which means that your reporting period must be within the length of time that your database is configured to store raw data. By default, DataCollector is configured to store raw data for 48 hours, but this value is [configurable](#).

Host-Pair Table — For the specific peer you selected on the peer table, this table displays information about each flow that the host received from the peer during the reporting period. Note that you will only be able to see flow data if your reporting period is within the window of time that raw data is stored in your database (48 hours by default). Additionally, because flows could have occurred at any time during the reporting period, each screen displays intervals of data and you will only see flow data if an actual flow occurred during the interval. To see the next interval, click **Next**. The table displays the following information for each flow:

Statistic	Description
Start Time	Date and time the first packet in the flow was seen (SysUpTime when first packet seen)
End Time	Date and time the last packet in the flow was seen (SysUpTime when last packet seen)
Duration (sec)	The number of seconds between the start time and the end time. For non-TCP flows, flow records are generally created one hour after PacketWise sees the last packet for the flow. Exceptions are transactional non-TCP flows, such as a DNS lookup over UDP or an ICMP ping. For these types of flows, the flow record is created when the transaction is completed.
Throughput (bps)	Rate of the flow (in bits per second)
Src Ip	IP address of the device that sent the flow
Src Port	Port on which the flow was sent
Dest Ip	IP address of the destination device
Dest Port	Port that received the flow
Service	Service associated with the traffic flow
Traffic Class	Inbound and outbound traffic class for the traffic flow
Efficiency	The percentage of bytes that did not require retransmission

Listener: 10.9.59.87 with Peer: 216.38.172.83

For View: 'Network'

For Traffic Class: /Inbound/NewYork/HTTP

From Tue., Dec 1, 2009 01:00:00 PM, GMT-0800 to Tue., Dec 1, 2009 02:59:59 PM, GMT-0800

[Previous](#)[Next](#)

Start Time	End Time	Duration	Throughput	Src Ip	Src Port	Dest Ip	Dest Port	Service	Traffic Class	Efficiency
2009-12-01 13:02:25-0800	2009-12-01 13:05:46-0800	201	5778094	216.38.172.83	80	10.9.59.87	64280	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:05:41-0800	2009-12-01 13:05:46-0800	5	128	216.38.172.83	80	10.9.59.87	61125	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:02:25-0800	2009-12-01 13:07:30-0800	305	16	216.38.172.83	80	10.9.59.87	57445	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:07:24-0800	2009-12-01 13:12:29-0800	305	16	216.38.172.83	80	10.9.59.87	61951	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:07:24-0800	2009-12-01 13:12:30-0800	306	23	216.38.172.83	80	10.9.59.87	65285	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:12:24-0800	2009-12-01 13:17:29-0800	305	16	216.38.172.83	80	10.9.59.87	63226	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:12:24-0800	2009-12-01 13:17:29-0800	305	23	216.38.172.83	80	10.9.59.87	54440	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:17:24-0800	2009-12-01 13:22:29-0800	305	16	216.38.172.83	80	10.9.59.87	54621	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:17:24-0800	2009-12-01 13:22:29-0800	305	15	216.38.172.83	80	10.9.59.87	54143	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:22:24-0800	2009-12-01 13:27:29-0800	305	16	216.38.172.83	80	10.9.59.87	58685	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:22:24-0800	2009-12-01 13:27:29-0800	305	16	216.38.172.83	80	10.9.59.87	63465	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:27:24-0800	2009-12-01 13:32:29-0800	305	16	216.38.172.83	80	10.9.59.87	51694	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 13:27:24-0800	2009-12-01 13:32:30-0800	306	15	216.38.172.83	80	10.9.59.87	65151	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 14:07:22-0800	2009-12-01 14:12:27-0800	305	16	216.38.172.83	80	10.9.59.87	63407	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 14:07:22-0800	2009-12-01 14:12:27-0800	305	15	216.38.172.83	80	10.9.59.87	51370	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 14:12:22-0800	2009-12-01 14:17:27-0800	305	17	216.38.172.83	80	10.9.59.87	55501	HTTP	/Inbound/NewYork/HTTP	100.00 %
2009-12-01 14:12:23-0800	2009-12-01 14:17:28-0800	305	16	216.38.172.83	80	10.9.59.87	61979	HTTP	/Inbound/NewYork/HTTP	100.00 %

[Previous](#)[Next](#)

Page: 1


Created on: Dec 2, 2009 4:54 PM

Drill Down: Top Talkers/Listeners for Class

Top 10 Pie Chart — Shows the bandwidth consumption of each top host relative to the others. Each pie slice represents the percentage of bandwidth the host consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates each host's IP address with a pie slice color.

Utilization Table — For traffic in the selected class, this table shows the following information for each of the top hosts:

Statistic	Description
Percent	The percentage of total bandwidth consumed by the host relative to the other top hosts in the selected class during the reporting period
Bytes	The total number of bytes used by the host
Flows	The total number of flows generated by the host
Packets	The total number of packets in the flows generated by the host
Peers	Number of peers the host has paired with
Efficiency	The percentage of bytes that did not require retransmission
Avg Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Avg Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the host detected during the report period



Clicking on an IP address allows you to [drill down](#) to details about the specific host pairs for the top talker/listener.

Drill Down: Top Talker/Listener Peers

To drill down to details about the specific peers for a talker or listener, you can click on an IP address in the host table. This drill-down option is available in all [Top Talker](#) and [Top Listener](#) reports (including [Top Listeners by Site](#) and [Top Talkers by Site](#)).

Top Talker/Listener Table — For the specific top talker/listener you selected at the top level of the report, this table displays a list of the peers from which the top host received traffic during the reporting period. For each peer entry, the table displays statistics for data that the top host received from each peer.

Top Listener: 10.9.59.87							
For View: 'Network'							
For Traffic Class: /Inbound/NewYork/HTTP							
From Sun., Nov 29, 2009 00:00:00 AM, GMT-0800 to Wed., Dec 2, 2009 03:00:00 PM, GMT-0800							
Peers	Bytes	Flows	Packets	Efficiency	Avg Total Delay	Avg Server Delay	Transactions
216.38.172.83	144.98 M	32	96.76 K	100.00	0.00	0.00	0
10.78.1.241	37.43 M	49.57 K	248.98 K	100.00	0.00	0.00	0
10.2.2.101	7.1 M	49.96 K	49.96 K	100.00	0.00	0.00	0
10.2.2.100	7.06 M	49.65 K	49.65 K	100.00	0.00	0.00	0
10.78.1.250	1.94 M	24.88 K	38.45 K	100.00	0.00	0.00	0
216.52.23.2	38.61 K	164	508	100.00	0.00	0.00	0
10.9.59.92	30.06 K	87	347	100.00	0.00	0.00	0
216.38.160.83	14.17 K	28	88	100.00	0.00	0.00	0
216.52.23.18	10.86 K	8	24	100.00	0.00	0.00	0
208.115.138.201	10.73 K	28	84	100.00	0.00	0.00	0
119.27.62.201	10.39 K	27	82	100.00	0.00	0.00	0
208.115.138.203	10.35 K	27	81	100.00	0.00	0.00	0
8.21.4.201	10.35 K	27	81	100.00	0.00	0.00	0
8.21.4.203	10.35 K	27	81	100.00	0.00	0.00	0
85.92.222.201	9.65 K	26	77	100.00	0.00	0.00	0
85.92.222.203	6.98 K	19	56	100.00	0.00	0.00	0
64.181.192.120	6.48 K	33	77	100.00	0.00	0.00	0
216.38.163.83	4.02 K	13	28	100.00	0.00	0.00	0
216.38.164.83	3.36 K	8	30	100.00	0.00	0.00	0
203.12.2.160	2.41 K	5	23	100.00	0.00	0.00	0
216.38.162.83	240	6	6	100.00	0.00	0.00	0

Dec 2, 2009 4:47 PM



For each peer, this table shows the following information:

Statistic	Description
Peer	The IP address of the peer
Bytes	The total number of bytes used by the host
Flows	The total number of flows generated by the host
Packets	The total number of packets in the flows generated by the host
Efficiency	The percentage of bytes that did not require retransmission
Avg Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Avg Server Delay	Average number of milliseconds required for servers to process requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.


Statistic	Description
Transactions	The total number of unique iterations of a transfer of data for the host detected during the report period

Clicking on a peer IP address or host name allows you to [drill down](#) to flow details for the specific listener-talker combination.

Drill Down: Top Talkers/Listeners for Application

To drill down to a list of top talkers or top listeners for the application, click on the  (talkers) or  (listeners) next to the application name. Note that this level of the report is based on FDR data, you will not be able to drill down to this report level.

Top Talker/Listener Chart and Table — For the specific application you selected at the top level of the report, this table displays a pie chart and list of top talkers/listeners during the reporting period. For each top host, the table displays statistics related to the selected application.



For the selected application's traffic during the report period, this table shows the following information for each top host:

Statistic	Description
Percent	The percentage of total bandwidth consumed by the host relative to the other hosts
Bytes	The total number of bytes used by the host
Flows	The total number of flows generated by the host
Packets	The total number of packets in the flows generated by the host
Peers	Number of hosts with whom the talker/listener has exchanged data
Efficiency	The percentage of bytes that did not require retransmission
Avg Total Delay	Average number of milliseconds to complete transactions; includes network delay and server delay
Avg Server	Average number of milliseconds required for servers to process

Statistic	Description
Delay	requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response.
Transactions	The total number of unique iterations of a transfer of data for the host detected during the report period

To drill-down to a list of peers for the selected talker or listener within the application, click on the IP address or hostname of the talker or listener.

Drill-Down: Peers for Listener/Talker

Top Talker/Listener Peers Table — For the specific talker or listener you selected, this table displays a list of peers with which the talker/listener exchanged flows for the selected application during the reporting period. For each peer, the table displays statistics related to the selected application (see the table above for a description of each column).

Top Listener: 10.9.50.93									
For Device Group: Network									
For Application: HTTP									
Special Report									
From Thu, Nov 20, 2008 01:00:00 PM, GMT-08:00 to Thu, Nov 20, 2008 01:59:59 PM, GMT-08:00									
Peers	Bytes	Flows	Packets	Efficiency	Avg Total Delay	Avg Server Delay	Transactions	PET Server	PET Client
10.9.50.92	117.96 K	5	106	100 %	0	0	0	0	0
69.147.64.75	29.29 K	6	45	100 %	0	0	0	0	0
68.180.219.132	14.58 K	34	168	100 %	0	0	0	0	848
76.13.210.11	6.72 K	1	14	100 %	0	0	0	0	0
216.73.87.152	2.51 K	1	5	100 %	0	0	0	0	0
216.252.124.207	2.05 K	3	15	100 %	0	0	0	0	0
66.94.234.72	627	1	5	100 %	0	0	0	0	0
88.136.43.76	607	1	5	100 %	0	0	0	0	0

Clicking the IP address of a peer allows you to drill down to the actual flows for the peer.

Drill-Down: Flow Details

Top Talker/Listener Peers Flows — Because this drill down shows actual flows, you will only see data for periods during which flows were exchanged between the selected peer and the specified talker/listener. The first time you drill down on a peer, you must specify the size of the interval IC uses to display the flow data: 5 minutes, 15 minutes, 30 minutes (default), 1 hour, or 2 hours. IC will then display the flow data for the selected interval, which you can scroll through using the **Next** or **Previous** links to view all flows within the reporting period. Note that flow data is stored as raw data and it can therefore only be displayed if it is still present in your database, which means that your reporting period must be within the length of time that your database is configured to store raw data. By default, DataCollector is configured to store raw data for 48 hours, but this value is [configurable](#).

Host-Pair Table — For the specific peer you selected on the peer table, this table displays information about each flow that the talker/listener exchanged with the peer during the reporting period. Note that you will only be able to see flow data if your reporting period is within the window of time that raw data is stored in your database (48 hours by default). Additionally, because flows could have occurred at any time during the reporting period, each screen displays intervals of data and you will only see flow data if an actual flow occurred during the interval. To see the next interval, click **Next**.

Start Time	End Time	Duration (sec)	Throughput (Bps)	Src IP	Src Port	Dest IP	Dest Port	Protocol	Service	OutboundPala_AkshHTTP	Peer	Ellipsis
Nov 20, 2008 12:58 PM	Nov 20, 2008 1:01 PM	108	40	88.168.219.137	80	10.9.58.93	1368	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 12:59 PM	Nov 20, 2008 1:01 PM	108	69	10.9.58.93	1368	88.168.219.137	80	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:00 PM	Nov 20, 2008 1:00 PM	6	1345	10.9.58.93	1371	88.168.219.137	80	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:00 PM	Nov 20, 2008 1:00 PM	6	525	88.168.219.137	80	10.9.58.93	1371	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:00 PM	Nov 20, 2008 1:01 PM	9	433	88.168.219.137	80	10.9.58.93	1373	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:00 PM	Nov 20, 2008 1:01 PM	9	899	10.9.58.93	1373	88.168.219.137	80	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:00 PM	Nov 20, 2008 1:04 PM	189	17	88.168.219.137	80	10.9.58.93	1377	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:00 PM	Nov 20, 2008 1:04 PM	189	38	10.9.58.93	1377	88.168.219.137	80	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:01 PM	Nov 20, 2008 1:01 PM	6	525	88.168.219.137	80	10.9.58.93	1379	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:01 PM	Nov 20, 2008 1:01 PM	6	1345	10.9.58.93	1379	88.168.219.137	80	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:02 PM	Nov 20, 2008 1:02 PM	6	1345	10.9.58.93	1380	88.168.219.137	80	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:02 PM	Nov 20, 2008 1:02 PM	6	925	88.168.219.137	80	10.9.58.93	1380	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:03 PM	Nov 20, 2008 1:03 PM	6	525	88.168.219.137	80	10.9.58.93	1381	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:03 PM	Nov 20, 2008 1:03 PM	6	1345	10.9.58.93	1381	88.168.219.137	80	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:04 PM	Nov 20, 2008 1:07 PM	178	27	88.168.219.137	80	10.9.58.93	1385	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:04 PM	Nov 20, 2008 1:07 PM	178	42	10.9.58.93	1385	88.168.219.137	80	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C
Nov 20, 2008 1:04 PM	Nov 20, 2008 1:04 PM	6	1345	10.9.58.93	1387	88.168.219.137	80	TCP	HTTP	/OutboundPala_AkshHTTP	3	10C

The table displays the following information for each flow:

Statistic	Description
Start Time	Date and time the first packet in the flow was seen (SysUpTime when first packet seen)
End Time	Date and time the last packet in the flow was seen (SysUpTime when last packet seen)
Duration (sec)	The number of seconds between the start time and the end time. For non-TCP flows, flow records are generally created one hour after PacketWise sees the last packet for the flow. Exceptions are transactional non-TCP flows, such as a DNS lookup over UDP or an ICMP ping. For these types of flows, the flow record is created when the transaction is completed.

Statistic	Description
Throughput (bps)	Rate of the flow (in bits per second)
Src IP	IP address of the device that sent the flow
Src Port	Port on which the flow was sent
Dest IP	IP address of the destination device
Dest Port	Port that received the flow
Protocol	IP protocol of the flow
Service	Service associated with the traffic flow
Class	Inbound and outbound traffic class for the traffic flow
Policy	Priority for this flow (0-7), either the priority assigned by a priority policy, or the priority assigned to excess rate with a rate policy
Efficiency	The percentage of bytes that did not require retransmission

Drill Down: Traffic Class Activity by Device or Group

In the [Top N Traffic Classes Summary](#) and [Top N Traffic Classes by Site](#) reports, you can click on a class name to drill down to details about the specific devices and groups responsible for generating the traffic that is classified into the selected traffic class.


Traffic Class Activity Table — For the specific traffic class you selected at the top level of the report, this table displays a list of the devices and/or groups that generated traffic for the class during the reporting period. For each device or group entry, the table displays the total number of packets and bytes used by the device or group, the partition size and utilization, the efficiency level percentage, and the average and peak rates for the traffic generated by the device or group.

Traffic Class Activity								
For Traffic Class: 'Inbound/Olympia'								
From Mon, Oct 04, 2010 12:00:00 AM GMT-0700 to Mon, Oct 11, 2010 12:00:00 AM GMT-0700								
IP Address	Device Name	Bytes	Packets	Partition Size	Partition Util (%)	Efficiency	Avg Rate (bps)	Peak Rate (bps)
10.78.52.51	10.78.52.51	1.37 T	1.55 G	0	0.00 %	98.46 %	34.56 M	163.68 M
10.78.52.49	10.78.52.49	1.23 T	1.4 G	0	0.00 %	98.52 %	33.34 M	149.86 M
10.78.52.50	10.78.52.50	1.23 T	1.4 G	0	0.00 %	98.41 %	33.13 M	156.29 M
Group	Bytes	Packets	Partition Size	Partition Util (%)	Efficiency	Avg Rate (bps)	Peak Rate (bps)	
Network	3.83 T	4.35 G	0	0.00 %	98.46 %	33.69 M	163.68 M	
Page: 1	Created on: Oct 11, 2010 10:44 PM							

Clicking on a specific device name or IP address or group name allows you to [drill down](#) to the Average Rate - Peak Rate Comparison for the device or group.

Drill Down: Traffic Class Rates Comparison

Traffic Class Utilization - Rates Comparison — Shows a line chart that details the average rate versus the peak rate for the specific traffic class-device or group combination you selected at the previous level of the report.



Portlets

IntelligenceCenter portlets are [customizable](#) applications that allow you to define your own unique views of your network. IntelligenceCenter provides the following portlets:

- [Application Performance Portlet](#)
- [Class Utilization Portlet](#)
- [Network Efficiency Portlet](#)
- [Per Subnet FDR Portlet](#)
- [Per Server FDR Portlet](#)
- [Top N Children Portlet](#)
- [VoIP Performance Portlet](#)

Application Performance Portlet

The Application Performance portlet allows you to monitor the health, service level agreement (SLA) compliance, and utilization of up to 10 TCP applications of interest. The Application Performance portlet gives information about application performance based on three network quality parameters: loss, latency, and availability. Keep in mind that the Application Performance portlet uses ME variables only in calculating application performance; therefore, if your DataCollector is not configured to [collect ME data](#), this portlet will be empty. The following table describes each quality parameter.

Parameter	Description
Loss	How much data is lost in the TCP flows of the application. High loss might be a result of poor WAN line quality, overutilization of network devices between the client and server, Ethernet collisions, or other situations where established TCP connections will lose data packets during the flow. Loss is calculated using the following ME variable formula: <code>tcp_retx_bytes / bytes</code>
Latency	A measure of the average packet exchange time (PET) for the TCP flows of the application, measured in milliseconds (msecs). High latency can be a result of sharing limited bandwidth with other applications, slow WAN links, or servers with high CPU utilization. The <code>pkt_exchange_time_msec</code> ME variable represents latency in the Application Performance calculation.
Availability	A measure of the availability of the network and server in accepting new TCP connections. Low availability could be a result of a server that can no longer accept new TCP connections for that application, power cycling servers or network equipment, or disconnected network devices between the client and server. Availability is calculated using the following ME variable formula: <code>tcp_conn_server_refuse + tcp_conn_server_ignores / tcp_conn_inits</code>

When you [configure the Application Performance portlet](#), you define baseline values for loss, latency, and availability. Because the degree to which each of these parameters is important depends on the application itself, you must also assign a weight to each parameter.

IntelligenceCenter computes the percentage difference between the baseline values you define and the actual values as follows:

$$(\text{baseline_value} - \text{actual_value} / \text{baseline_value}) * 100\% = \% \text{difference}$$


Note: If the actual value is better than the baseline value, the %difference is set to 0; if the %difference is higher than 100%, it is set to 100%.

It then uses the %difference values along with the weights you assigned to each parameter to determine the application's level of compliance for a given monitoring period as follows:


$$100\% - (\% \text{difference_loss} * \text{weight_loss} + \% \text{difference_latency} * \text{weight_latency} + \% \text{difference_availability} * \text{weight_availability}) = \text{application performance}$$

IntelligenceCenter then compares the application performance value for each application to the baseline values you specified for Normal (green), Warning (orange), and Critical (red) and displays the results in the portlet window in three graphs:


Performance Summary bar graph — Shows the application performance of each monitored application during the specified time period. Each section of a bar represents the proportion of time during the specified date range that the application is at the Critical (performance-violation), Warning, or Normal (performance-compliant) level. For example, if an application shows 78% Normal, this means that the application was in compliance with the performance threshold 78% of the days in the date range. Hover over a bar to see specific percentages.



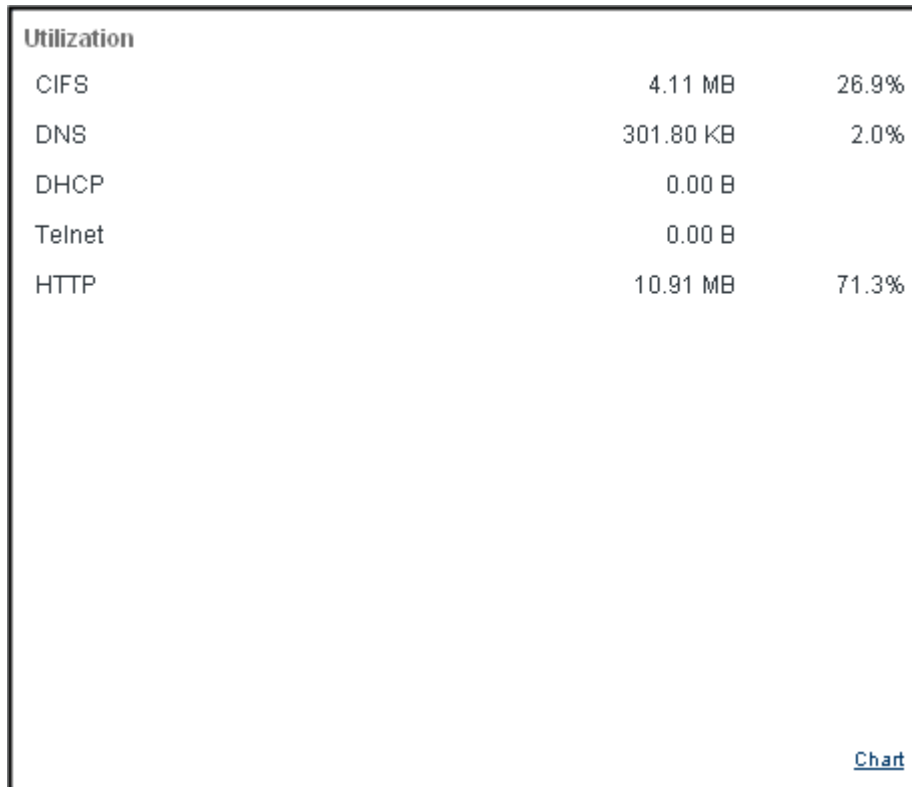
Detail Information bar graph — When you click an application's bar in the *Performance Summary*, another graph appears in the *Detail* pane. Each bar represents the health of the application for a particular date or time. The health percentage is based on the application's availability, latency, and loss; it takes into account the baselines and weights assigned to the application definition. Hover over a bar to see specific percentages.



Utilization pie chart — Shows the bandwidth utilization of each monitored application for the specified time period. The value of each pie slice is indicated on the chart (MB=Megabytes, GB=Gigabytes, TB=Terabytes). Hover over a slice to see specific percentages and number of bytes.



To view utilization statistics for each application, click the **Data** link below the pie chart. You can toggle back to the pie chart by clicking the **Chart** link.



By default, the Application Performance portlet measures data across the entire network. However, you can [edit the portlet settings](#) to specify which network branch or view you want to monitor. You can also customize which applications you want to monitor.

Note: Before you can monitor an application using the Application Performance portlet, IntelligenceCenter must know about the application. IntelligenceCenter automatically creates applications based on the traffic classes defined on your network. However, if you have custom traffic classes or if you want to monitor non-standard applications, you will need to [create your own application definitions](#) or [modify the application definitions](#) that IntelligenceCenter dynamically creates.

Keep in mind that the time range you select when customizing the Application Performance portlet determines what [database table](#) DataCollector will pull the data from. Therefore, when selecting a time range, you must consider whether DataCollector has had a chance to roll data up into the corresponding database table. If the corresponding database table does not contain any data yet, the portlet will not display any data. For example, if you select a seven day time range, IntelligenceCenter will pull the data from the DataCollector's daily table. However, if you do not yet have any data rolled up into the daily table, the portlet cannot display any data.

When you first add a portlet to a portlet view, the default time range for the portlet is This week, which means it will get its data from the daily table. However, if you have not yet collected a day's worth of data the daily table will not contain any data and the portlet will therefore not display any data. If your portlet does not display data, try selecting a shorter time range.

Class Utilization Portlet

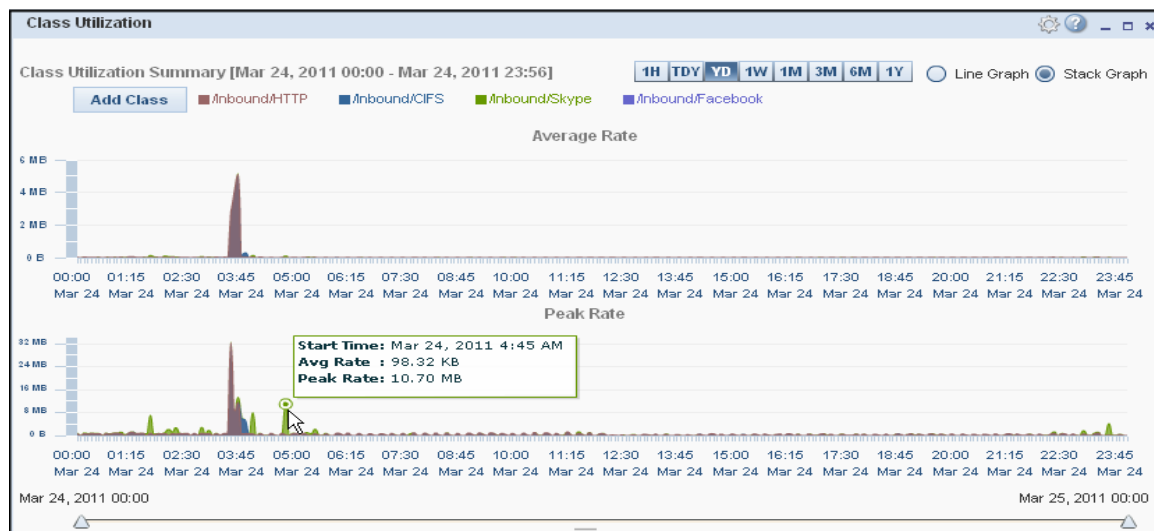
The Class Utilization portlet provides at-a-glance monitoring of the traffic on your network. With each instance of the Class Utilization portlet that you configure, you can monitor and compare the average and peak bandwidth usage of up to five traffic classes. This allows you to group classes for specific monitoring purposes.

For example, application administrators might want to configure the portlet to monitor traffic classes that correspond to the enterprise's mission critical applications, enabling them to take action immediately when the performance of a critical class starts to degrade.

Network administrators, on the other hand, might want to create portlet instances to help them troubleshoot an issue that has been reported on a particular PacketShaper or at a specific site. In this case, they can leverage the portlet's bandwidth monitoring capabilities to look for signs of network health issues, such as flapping utilization values, which may indicate excessive TCP aborts or time-outs. Once the issue is identified, they can then use the portlet to aid in troubleshooting by viewing average and peak bandwidth usage over various periods of time. For example, by zooming out the administrator can identify when the issue first began and get an idea of its frequency and duration. Similarly, by zooming in on the data, the administrator can pinpoint the exact time of an issue to aid in further troubleshooting.


This portlet is based on ME data. If you are not collecting ME data you will not be able to use this report. Also, because there is a delay in the collection of raw ME data, the portlet only shows data up through the last full quarter of an hour.



After you [configure the Class Utilization portlet](#), you can begin monitoring the traffic classes. The top part of the portlet graphs the selected traffic classes average and/or peak bandwidth consumption, in bits per second, over the reporting period. To see the values at a specific point in time, hover over the line or area. By default, the portlet shows the graphs as line graphs. You can toggle the type of graph that the portlet displays by selecting the **Stack Graph** and **Line Graph** radio buttons.

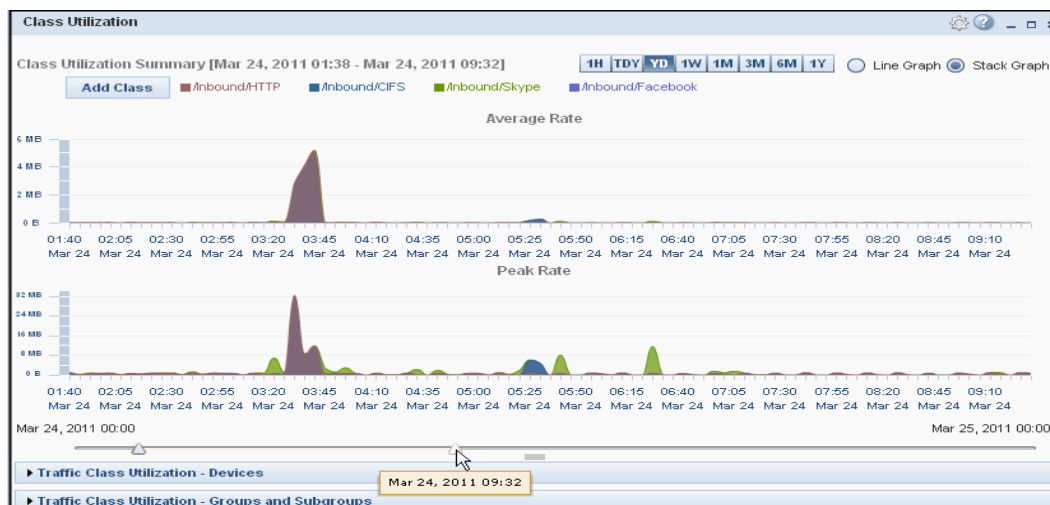


- Click the ► and ▼ icons to show and hide the Traffic Class Utilization tables and the class data within each table. If you want more room for viewing the tables. These tables display detailed utilization information for each of the classes you are monitoring, including number of bytes and packets, partition size and utilization, efficiency (expressed as a percentage of bytes that did not require retransmission), and the average and peak rates (in bits per second). Utilization details are broken down by the devices, groups and subgroups that reported the class data that is currently aggregated on the graphs.

▼ Traffic Class Utilization - Devices									
IP Address	Device Name	Bytes	Packets	Partition	Partition Util	Efficiency	Avg Rate	Peak Rate (bps)	
▼ Network:Inbound/CIFS									
10.125.32.20	117-10014614	16.31 MB	283.84 KB	0.00 B	0.00%	100.00%	1.52 KB	5.27 MB	
▼ Network:Inbound/Facebook									
10.125.32.20	117-10014614	102.78 KB	102.00 B	0.00 B	0.00%	96.00%	9.00 B	80.26 KB	
▼ Network:Inbound/HTTP									
10.125.32.20	117-10014614	584.59 MB	496.09 KB	1.00 MB	5.00%	99.00%	54.22 KB	32.29 MB	
▼ Traffic Class Utilization - Groups and Subgroups									
Group	Bytes	Packets	Partition Size	Partition Util (%)	Efficiency	Avg Rate (bps)	Peak Rate (bps)		
▼ Network:Inbound/CIFS									
BVSubGroup	16.31 MB	283.84 KB	0.00 B	0.00%	100.00%	1.52 KB	5.27 MB		
Network	16.31 MB	283.84 KB	0.00 B	0.00%	100.00%	1.52 KB	5.27 MB		
▼ Network:Inbound/Faceb									
BVSubGroup	102.78 KB	102.00 B	0.00 B	0.00%	96.00%	9.00 B	80.26 KB		
Network	102.78 KB	102.00 B	0.00 B	0.00%	96.00%	9.00 B	80.26 KB		

Tip: You can click and drag the  to move the divider between the graphs and the tables if you need more room for one or the other.

- To ensure that you are viewing the most up to date information, click the  icon in the portlet's title bar and select **Refresh Data** from the pop-up menu.
- If you see an interesting area on the graph—for instance a period in which the peak rate for one of your monitored classes spikes or a period where the average rates drop off—you can zoom in on that area of the graph using the  sliders below the graph. Moving the slider on the left side of the graph moves the start time of the reporting period; moving the slider on the right side of the graph moves the end time of the reporting period. Before moving the sliders, note the time where the part of the graph you are interested in occurs; the time is marked on the x-axis of each graph. As you move the sliders, a tool tip displays the time you are moving to help you pinpoint where to stop dragging. As you slide, the graph redraws with the new start or end time.



- You can modify the set of traffic classes the portlet is monitoring as follows:

- To remove a traffic class, hover over the traffic class name in the legend and then click the **Remove Class** icon.




- To add a traffic class, click **Add Class**. See "Configuring the Class Utilization Portlet" above for detailed instructions on how to add a traffic class. You can add up to five traffic classes to each Class Utilization portlet instance.


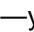



Network Efficiency Portlet

The Network Efficiency portlet allows you to monitor the efficiency of the TCP traffic classes on your network. For each traffic class that you select, the portlet graphs the network efficiency percentage, which represents the percentage of bytes that are *not* retransmits. Tosses and retransmitted packets lower the efficiency percentage. An efficient network, displayed as minor fluctuations from 100%, needs very little intervention. A lower percentage indicates that more of your network capacity is devoted to retransmitting packets, a sign of congested router queues, delays, and time-outs. In this case, you may want to reevaluate your current partitions and policies.

After you [configure the Network Efficiency portlet](#), you can begin monitoring the efficiency of the traffic classes you selected:



- Hover over a data point to see details.
- To ensure that you are viewing the most up to date information, click the  icon in the portlet's title bar and select **Refresh Data** from the pop-up menu.
- If you see an interesting area on the graph—for instance a period in which the efficiency level drops—you can zoom in on that area of the graph using the  sliders below the graph. Moving the slider on the left side of the graph moves the start time of the reporting period; moving the slider on the right side of the graph moves the end time of the reporting period. Before moving the sliders, note the time where the part of the graph you are interested in occurs; the time is marked on the x-axis of each graph. As you move the sliders, a tool tip displays the time you are moving to help you pinpoint where to stop dragging. As you slide, the graph redraws with the new start or end time.




You can modify the set of traffic classes the portlet is monitoring as follows:

- To remove a traffic class, hover over the traffic class name in the legend and then click the **Remove Class** icon.



- To add a traffic class, click **Add Class**. See "Configuring the Class Utilization Portlet" above for detailed instructions on how to add a traffic class. You can add up to five traffic classes to each Class Utilization portlet instance.



Per Server FDR Portlet

The Per Server FDR portlet tracks recent traffic flows to your inside servers and allows you to keep an eye on traffic that you consider to be suspect. When you [select a server](#), this portlet provides information about the number of flows of for each service and allows you to drill down to actual flow detail record (FDR) information.

Per Server FDR

Group Network

Server: 255.255.255.255 Last 30 minutes Refresh

End Time 12:00 PM Updated 12:01 PM

Summary

☐ Abnormal ☒ All

Service	Flows	Bytes	Packets
SLP	1	336.00 B	4
NetBIOS-IP	1	229.00 B	1

When you drill down to service flow details, a new tab opens that displays information about the host that generated the flow, including number of flows, bytes, and packets.

Per Server FDR

Group Network

Server: Last 30 minutes [Refresh](#)

End Time 12:00 PM Updated 12:01 PM

Summary **NetBIOS-IP**

Client	Flows	Bytes	Packets
10.125.32.86	3	536.00 B	3

If you determine that a service is normal and does not present a threat, you can leave it as is. If you determine that it does present a threat, click the corresponding icon to disapprove usage of the service.

Per Server FDR



Group Network

Server: Last 30 minutes [Refresh](#)

End Time 12:00 PM Updated 12:01 PM

Summary **NetBIOS-IP**

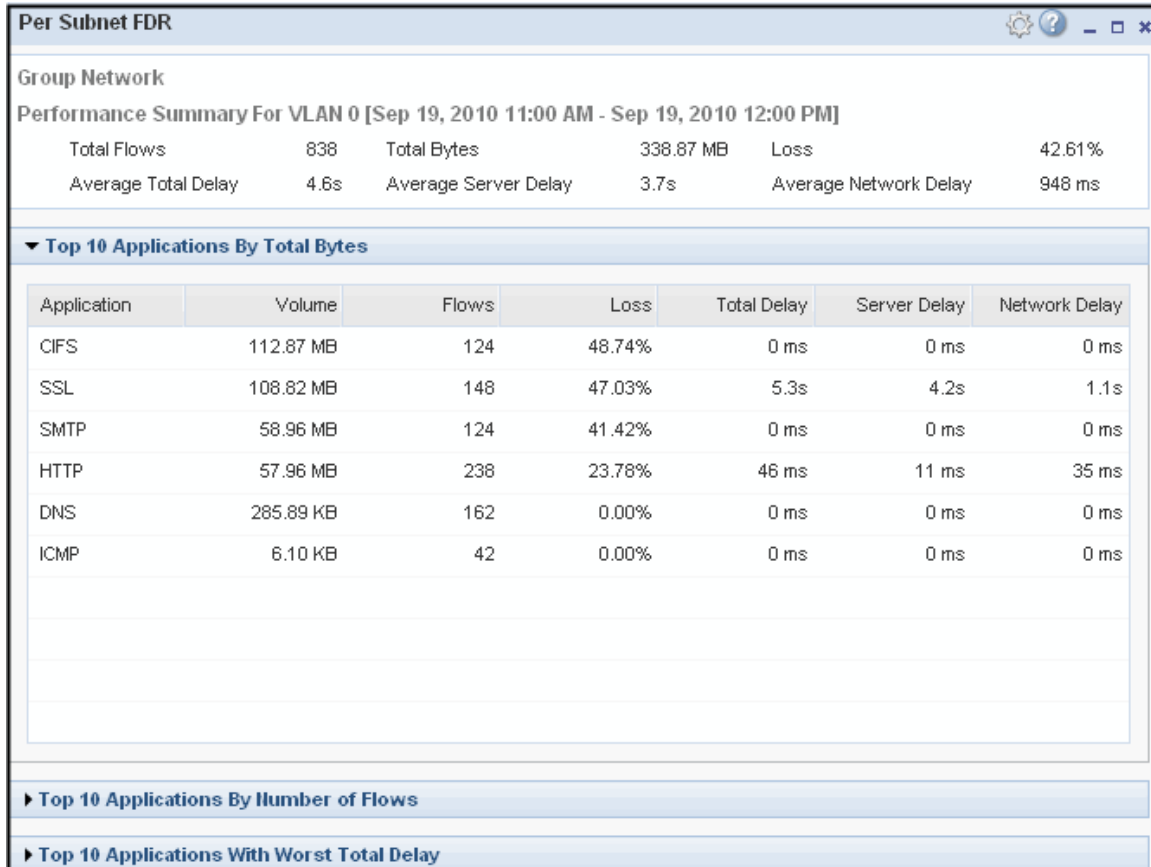
☐ Abnormal ☒ All

Service	Flows	Bytes	Packets
 SLP	<u>1</u>	336.00 B	4
 NetBIOS-IP	<u>1</u>	229.00 B	1

Click to approve the usage of this service

Per Subnet FDR Portlet

The Per Subnet FDR portlet displays statistics for traffic between two sites (such as two subnets) or on a specific virtual LAN (VLAN). This portlet allows you to monitor applications to see which ones are generating the most traffic in a specific time period. The Per Subnet FDR portlet is also useful for spotting application performance problems.



Note: The Per Subnet FDR portlet uses Packeteer-2 [Flow Detail Record](#) (FDR) data to calculate its values (that is, it does not use ME data or NetFlow-5 data). Therefore, if your DataCollector is not [configured to collect](#) Packeteer-2 data, this portlet will not display any data.

Summary Information

The *Summary* section at the top lists the following statistics for the specified location during the specified time period.

Field	Description
Total Bytes	Total number of bytes for all flows, measured in megabytes
Total Flows	Total number of flows
Loss (%)	Percentage packet loss The Loss field is calculated as follows: $\text{retransmitted_bytes} / \text{total_bytes}$
Total Delay (ms)	The total delay is the time required to complete a transaction; it includes network and server delay. The Total Delay field is calculated as follows: $\text{sum}(\text{total_delay}) / \text{sum}(\text{total_transactions})$ In other words, the delay is normalized by the number of transactions.
Server Delay (ms)	Server delay is the time required for servers to process transaction requests. The Server Delay field is calculated as follows: $\text{sum}(\text{server_delay}) / \text{sum}(\text{total_transactions})$
Network Delay (ms)	Network delay is the time a transaction spends in transit. The Network Delay field is calculated as follows: $\text{total_delay} - \text{server_delay}$

Application Statistical Tables

The Per Subnet FDR portlet includes three statistical tables that list the top ten applications by bytes, flows, or delay. Each of these tables is described below.

Top 10 Table	Description
Top 10 Applications by Total Bytes	Lists the applications that have the most traffic, measured in bytes, at the specified location during the specified time period. For each application, the table displays the total number of bytes, total number of flows, percentage packet loss, and response time measurement statistics (total delay, server delay, and network delay).
Top 10 Applications by Number of Flows	Lists the applications that have the most flows at the specified location during the specified time period. For each application, the table displays the total number of bytes, total number of flows, percentage packet loss, and response time measurement statistics (total delay, server delay, and network delay).
Top 10 Applications with Worst Total Delay	Lists the applications that have the worst response time (that is the most total delay) at the specified location during the specified time period. For each application, the table displays the total number of bytes, total number of flows, percentage packet loss, and response time measurement statistics (total delay, server delay, and network delay).

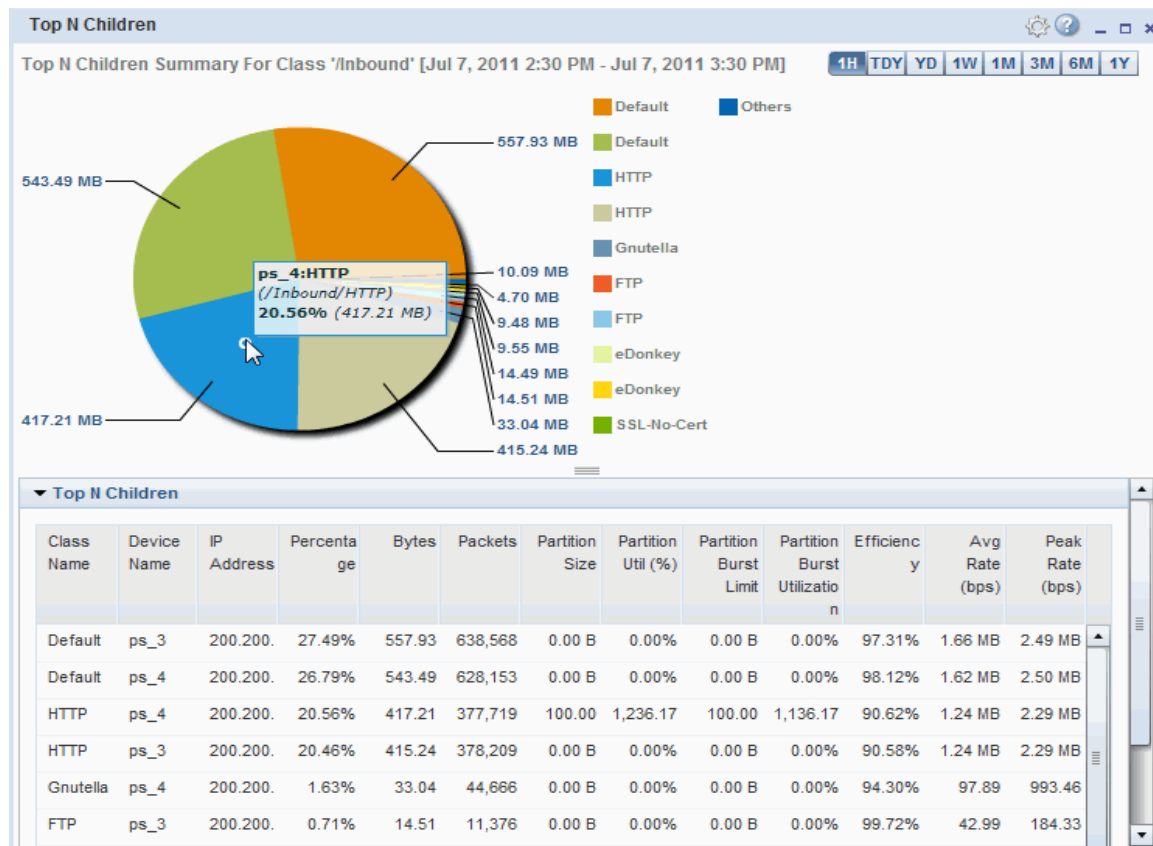
In order for an application to show up on the Per Subnet FDR application statistical tables, IntelligenceCenter must know about the application. IntelligenceCenter automatically creates applications based on the traffic classes defined on your network. However, if you have custom traffic classes or if you want to monitor non-standard applications, you will need to [create your own application definitions](#) or [modify the application definitions](#) that IntelligenceCenter dynamically creates.

Keep in mind that the time range you select when customizing the Per Subnet FDR portlet determines what [database table](#) DataCollector will pull the data from. Therefore, when selecting a time range, you must consider whether DataCollector has had a chance to roll data up into the corresponding database table. If the corresponding database table does not contain any data yet, the portlet will not display any data. For example, if you select a seven day time range, IntelligenceCenter will pull the data from the DataCollector's daily table. However, if you do not yet have any data rolled up into the daily table, the portlet cannot display any data.

Top N Children Portlet

The Top N Children portlet allows you to monitor the relative portions of bandwidth allocated to the ten most active children classes of the selected traffic class. This portlet is based on ME data. If you are not collecting ME data, you will not be able to generate this report.

After you [configure the Top N Children portlet](#), you can begin monitoring the bandwidth usage of the child classes for the selected traffic class:



This portlet contains two sections:

Pie Chart—Graphs each child class' average bandwidth usage in bits per second and its percentage of the parent class' total bandwidth usage. The pie chart shows the number of classes you specified when you [configured the portlet](#) plus an additional slice called *Others* that aggregates the bandwidth usage of all other child classes. Each pie slice represents the percentage of bandwidth the traffic class consumed during the reporting period. Hover over a slice to see specific percentages. The legend to the right of the pie chart associates a traffic class with a pie slice color.

Child Class Utilization Table—For each of the top children, this table shows the name of the class, the name and IP address of the PacketShaper appliance that reported the class data, the percentage of the parent class' total bandwidth usage, the total number of packets and bytes used by the class, the partition size and utilization as well as the partition burst limit and burst utilization, the efficiency percentage (ratio of bytes not requiring retransmission to the total number of bytes sent), and the average and peak rates for the class during the reporting period. You can sort the table by clicking the column header by which to sort. For example, click the **Percentage** column header to order the traffic classes according to bandwidth usage. When you sort the table by a particular column, an arrow appears next to the column header. An up arrow ▲

icon indicates that the data is being sorted in ascending order and a down arrow ▼ icon indicates that the data is being sorted in descending order. To switch the order in which the reports are sorted (ascending or descending), click the column header again.

You can also change the time range for which the portlet displays data by clicking one of the time span buttons at the top of the portlet.



VoIP Performance Portlet

The Voice over Internet Protocol (VoIP) Performance portlet allows you to monitor the performance of the RTP- and RTCP-based VoIP calls running on your network. The quality of the user experience for VoIP calls can be calculated based on three statistics:

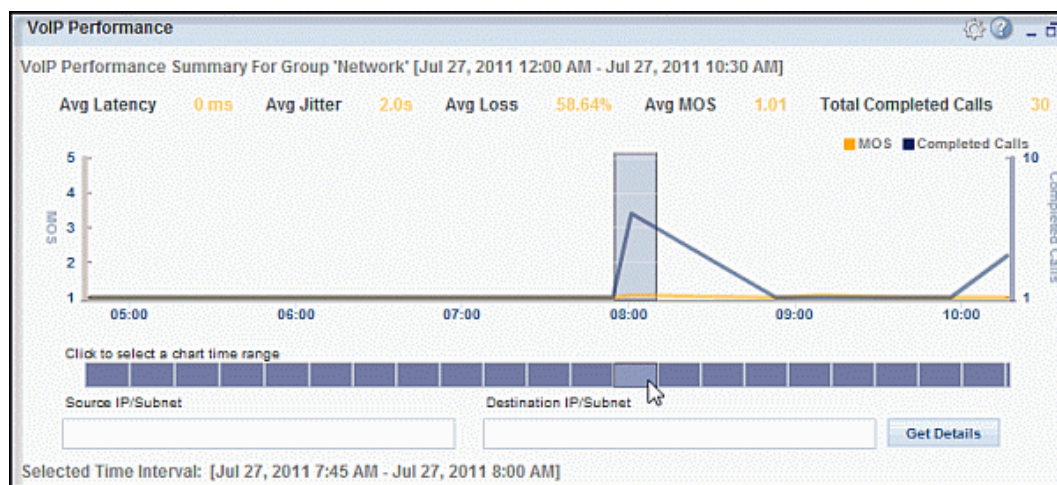
- **Loss**—the percentage of lost packets
- **Latency**—the time required for packets to travel from one PacketShaper to another. Latency is calculated on the Inbound interface only and can only be calculated between PacketShapers (one to intercept the inbound traffic at each end of a call). If you do not have a PacketShaper at each end of the VoIP traffic, zero will be used as the latency value.
- **Jitter**—the variation in the delay of received packets in a flow

Additionally, DC uses the statistics it collects to calculate a mean opinion score (MOS) for the VoIP traffic. The MOS value ranges from 1 (worst) to 5 (best) and provides a relative measure of VoIP quality.

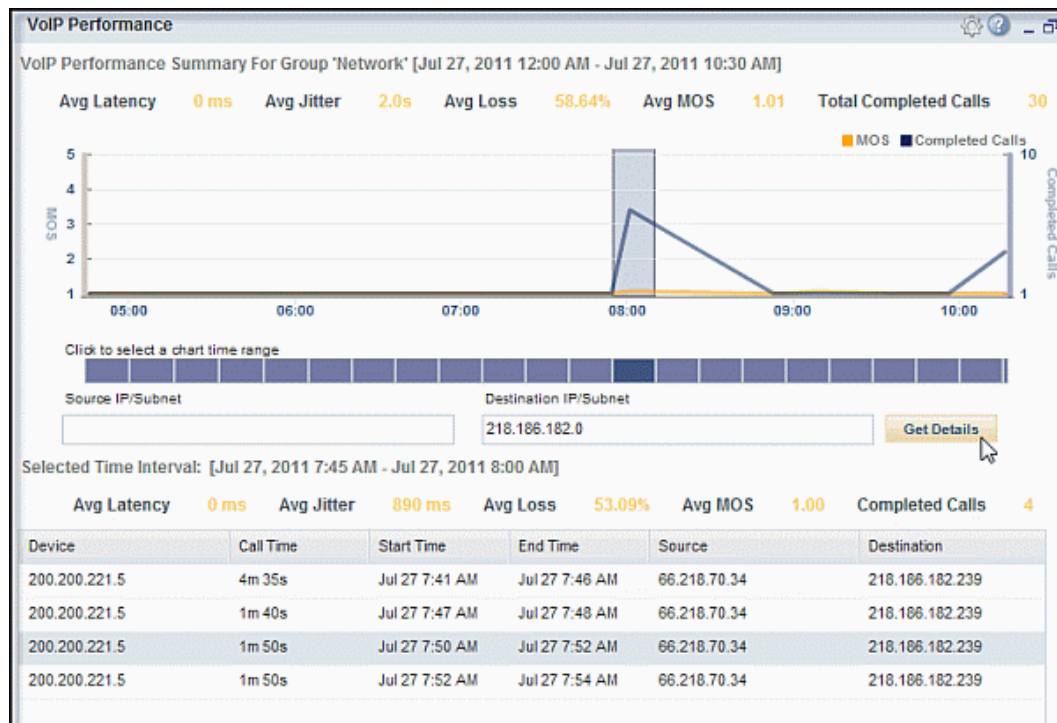
After you [configure the VoIP Performance portlet](#), you can begin monitoring the VoIP calls on your network.


The portlet displays summary statistics and a graph that show the average values for latency, jitter, loss, and MOS as well as total number of VoIP flows during the selected time span. You can drill down to see actual flow detail record (FDR) data that detail the VoIP calls for a specific 15-minute interval of the graph as follows:

1. On the time-selector bar below the graph, click on a section that corresponds to section of the graph for which you want to view call data. When you click a spot on the time-selector bar, the corresponding 15-minute interval on the graph is highlighted.



2. After you select the 15-minute interval for which you want to view call records, click **Get Details**. The portlet displays the average values for latency, loss, jitter, and MOS as well as total number of calls across the highlighted 15-minute interval as well as the FDR records for each VoIP flow (up to the maximum number you specified when you configured the portlet; 200 by default). Keep in mind that because VoIP traffic is UDP, each call record represents half of a VoIP call (caller to receiver or receiver to caller). For each record, the table shows the device that reported the flow, the start time, end time, and duration of the flow, the source and destination IP addresses of the flow, as well as the latency, jitter, loss, and MOS value for the flow. Also, if you defined your own latency value when you configured the portlet, the latency values on the table will display in blue.



3. If you want to filter the flow records to only show flows for a specific source and/or destination IP address or subnet, enter values in the **Source IP/Subnet** and/or **Destination IP/Subnet** field(s) and then click **Get Details**.
4. If you want to show flow records for a different 15-minute interval, select a different area on the time-selector bar and then click **Get Details**.
5. If you want to ensure that you are viewing the most up-to-date, click the  icon in the portlet's title bar and select **Refresh Data** from the pop-up menu.

Reference

This section contains the following reference items:

- [Measurement Variables](#) — Lists all of the common, link, partition, class, and compression variables that DataCollector collects
- [Report Options](#) — Lists the report options that you must configure when running reports
- Glossary — Lists common IntelligenceCenter terms

Measurement Variables

The measurement data that DataCollector collects depends on the appliance model. DataCollector collects measurement data from the following appliances:

- [PacketShaper](#)
- [PacketShaper ISP](#)

Measurement Variables Collected from PacketShaper

This section lists the measurement variables that DataCollector collects from PacketShaper (running PacketWise version 7.3.1 or higher) and PS-S500 appliances (running version 11.1.1.14 or higher). It collects a different set of variables from [PacketShaper ISP](#).

Each measurement type — link, partition, and class — also has its own set of variables. The common variables are available for all measurement types. The following tables list the variables by type: Common, Link, Partition, and Class.

In addition, DataCollector collects compression measurement variables, which are shown in their own table. These variables are specific to the PacketShaper compression feature. The standard measurement variables (link, partition, class) measure the actual data (compressed and uncompressed) that passes through the link. The compression variables allow you to analyze the effectiveness of compression on your link.

Common Variables

The common variables are available for link, partition, and class measurement.

Common Variables	Description
bytes	Number of bytes that passed through the unit during the specified interval
guar-rate-fails	Count of "denied" admission-control events when there wasn't enough bandwidth to satisfy guaranteed rate demand. This is an indicator of guaranteed rate failures and therefore can be used as a measurement of user satisfaction.
peak-bps	Peak rate recorded for the link, partition, or class during the specified interval; PacketWise determines the peak-bps by looking at the rate recorded for the busiest one-second subinterval (that is, the subinterval that had the highest rate).
pkts	Packet count — number of packets that passed through the unit during the specified interval
sample-interval-secs	Time between measurement samples (interval duration) in seconds
tcp-conn-aborts	Number of TCP connections exited as aborted connections; for example, HTTP stop button hits. This number gets incremented when a connection is in progress and an RST is sent. (RSTs are sometimes used as a lazy way to close a connection as an alternative to sending the full orderly shutdown FIN, FIN-ACK, ACK sequence to close a TCP connection.)
tcp-conn-inits	Number of TCP connections started. This number gets incremented when PacketWise sees a SYN packet initiating a new flow.
tcp-conn-server-ignores	Number of TCP connections exited as ignored connections — that is, the server never responded. This number gets incremented when a connection is quarantined (which means the flow limit specified in the flowlimit policy was exceeded) or if either the SYN ACK or SYN ACK ACK are not seen within a one-minute time-out period.
tcp-conn-server-refuses	Number of TCP connections refused by the server. This number gets incremented when a SYN is refused using an RST. Generally this happens when a server wants to deny a connection because it's too busy to accept a new connection.
tcp-retx-bytes	Count of TCP retransmitted bytes

Link Variables

DataCollector collects the following link variables.

Link Variables	Description
link-size-bps	Configured link size
total-rx-bytes	Total number of bytes received, not including bytes that were dropped
total-rx-packets	Total number of packets received, not including packets that were dropped
total-tx-bytes	Total number of bytes transmitted, not including bytes that were dropped
total-tx-packets	Total number of packets transmitted, not including packets that were dropped.

Partition Variables

DataCollector collects the following partition variables.

Partition Variables	Description
partition-burst-limit-bps	Configured partition burst limit; same as the partition size if the partition is not burstable
partition-over-limit-msecs	Cumulative time during which allocated bandwidth exceeded the partition minimum size, in milliseconds
partition-size-bps	Configured partition minimum size

Class Variables

DataCollector collects the following class variables.

Class Variables	Description
class-hits	Number of times flows match the class; class hits occur only at the beginning of a flow or session
is-rtp	Indicates whether the traffic class is VoIP traffic.
network-delay-msec	The sum of the network delays of all transactions in the specified interval, measured in milliseconds. (Network delay is the time a transaction spends in transit.) This variable is useful for calculating weighted averages across multiple intervals.
pkt-exchange-time-msecs	Interval between when a data packet leaves the PacketShaper and its ACK arrives (ACKs to SYN packets are not counted)
pkt-exchange-time-samples	Number of pkt-exchange-time samples that were counted
round-trip-time-msecs	The sum of the round-trip-times (RTT) of all transactions in the specified interval, measured in milliseconds. (RTT is the time a packet takes to go from client to server and back again.) Note that this measurement is taken once per transaction (not once per packet). This variable is useful for calculating weighted averages across multiple intervals.
rtp-expected-packets	Expected number of RTP packets
rtp-jitter-count	Number of jitter samples used to compute jitter
rtp-jitter-sum	Sum of measured jitter
rtp-latency-count	Number of RTP latency samples used to compute latency
rtp-latency-sum	Sum of measured RTP latency
rtp-lost-pkts	Number of lost RTP packets
server-delay-msec	The sum of the server delays of all transactions in the specified interval, measured in milliseconds. (Server delay is the time required for servers to process the class' transaction requests.) This variable is useful for calculating weighted averages across multiple intervals.
slow-transactions	Number of slow transactions (as defined by the total-delay-threshold variable)
total-delay-msec	The sum of the delays of all transactions in the specified interval, measured in milliseconds. (Total delay is the time required to

Class Variables	Description
	complete a transaction; includes network and server delay.) This variable is useful for calculating weighted averages across multiple intervals.
total-trans	Number of transactions (request-response pairs)
trans-bytes	Transaction size for TCP-based applications

Compression Variables

DataCollector collects the following compression variables. Keep in mind that the compression measurement variables, unlike the regular ME variables, do not include link and header overhead. Therefore, it does not make sense to compare the compression measurement variables with the regular link, class and partition variables.

Compression Variables	Description
compression-mode	Indicates whether compression is on, off, or has changed during the interval (1=off, 2=changed, 3=on) In migration mode, if either legacy compression or enhanced compression is on, the value of compression-mode will be 3 (on).
non-compressible-bytes	Number of bytes that PacketWise didn't attempt to compress either because they didn't belong to a compressible service or because they weren't compressible
postcompression-bytes	Number of bytes that actually passed through the unit. It includes bytes that were compressed as well as bytes that weren't compressible.
postcompression-peak-bps	Peak rate recorded for the link when compression is enabled (similar to peak-bps except it doesn't include header or link overhead); includes compressible and non-compressible traffic
precompression-bytes	Number of bytes passing through the unit before compression has been applied; includes compressible and non-compressible traffic

Measurement Variables Collected from PacketShaper ISP

This section lists the measurement variables that DataCollector collects from PacketShaper ISP appliances. DC can only collect data from PacketShaper ISP appliances that are running PacketWise version 7.3.1 or higher. It collects a different set of variables from [PacketShaper appliances](#).

Each measurement type — link, partition, and class — also has its own set of variables. The common variables are available for all measurement types. The following tables list the variables by type: Common, Link, Partition, and Class.

Common Variables

The common variables are available for link, partition, and class measurement.

Common Variables	Description
bytes	Number of bytes that passed through the unit during the specified interval
peak-bps	Peak rate recorded for the link, partition, or class during the specified interval; PacketWise determines the peak-bps by looking at the rate recorded for the busiest one-second subinterval (that is, the subinterval that had the highest rate).
pkts	Packet count — number of packets that passed through the unit during the specified interval
sample-interval-secs	Time between measurement samples (interval duration) in seconds
tcp-conn-aborts	Number of TCP connections exited as aborted connections; for example, HTTP stop button hits. This number gets incremented when a connection is in progress and an RST is sent. (RSTs are sometimes used as a lazy way to close a connection as an alternative to sending the full orderly shutdown FIN, FIN-ACK, ACK sequence to close a TCP connection.)
tcp-conn-inits	Number of TCP connections started. This number gets incremented when PacketWise sees a SYN packet initiating a new flow.
tcp-conn-server-ignores	Number of TCP connections exited as ignored connections — that is, the server never responded. This number gets incremented when a connection is quarantined (which means the flow limit specified in the flowlimit policy was exceeded) or if either the SYN ACK or SYN ACK ACK are not seen within a one-minute time-out period.
tcp-conn-server-refuses	Number of TCP connections refused by the server. This number gets incremented when a SYN is refused using an RST. Generally this happens when a server wants to deny a connection because it's too busy to accept a new connection.
tcp-retx-bytes	Count of TCP retransmitted bytes

Table of Link Variables

DataCollector collects the following link variables.

Link Variables	Description
link-size-bps	Configured link size
total-rx-bytes	Total number of bytes received, not including bytes that were dropped
total-rx-packets	Total number of packets received, not including packets that were dropped
total-tx-bytes	Total number of bytes transmitted, not including bytes that were dropped
total-tx-packets	Total number of packets transmitted, not including packets that were dropped.

Partition Variables

DataCollector collects the following partition variables.

Partition Variables	Description
partition-burst-limit-bps	Configured partition burst limit; same as the partition size if the partition is not burstable
partition-size-bps	Configured partition minimum size

Class Variables

DataCollector collects the following class variables.

Class Variables	Description
class-hits	Number of times flows match the class; class hits occur only at the beginning of a flow or session

CSV Record Formats

You can back up your FDR data to (comma-separated values) CSV files. The column headers for the back up files depend on the record format. You can back up the following FDR record formats to a CSV file:

- [Packeteer-2](#)
- [NetFlow-5](#)

Packeteer-2 Record Format

The following describes the columns in the Packeteer-2 CSV file:

Column Name	Description
Device ID	Serial number of the PacketShaper appliance
Flow ID	PacketShaper flow identifier
Source IP Address	IP address of the device that sent the flow
Destination IP Address	IP address of the destination device
Class ID	Traffic class into which the flow is classified, represented as an identification number
Inbound Interface	Index number that identifies the Inbound PacketShaper interface that packets from this flow initially arrived on. The interface mapping depends on the appliance type, but can be any of the following: 1 Inside (built-in) 2 Outside (built-in) 3 Upper_Inside (PacketShaper 2500, 6500, 8500, 9500, and 10000), Right-Inside (PacketShaper 3500 and 7500), or Backup_Inside (PacketShaper 1400) 4 Upper_Outside (PacketShaper 2500, 6500, 8500, 9500, and 10000), Right-Outside (PacketShaper 3500 and 7500), or Backup_Outside (PacketShaper 1400) 5 Lower_Inside (PacketShaper 2500, 6500, 8500, 9500, and 10000) or Left_Inside (PacketShaper 3500 and 7500) 6 Lower_Outside (PacketShaper 2500, 6500, 8500, 9500, and 10000) or Left_Outside (PacketShaper 3500 and 7500)
Outbound Interface	Index number that identifies the Outbound PacketShaper interface that packets from this flow left on. The interface mapping depends on the appliance type, but can be any of the following: 1 Inside (built-in) 2 Outside (built-in) 3 Upper_Inside (PacketShaper 2500, 6500, 8500, 9500, and 10000), Right-Inside (PacketShaper 3500 and 7500), or Backup_Inside (PacketShaper 1400)

Column Name	Description
	4 Upper_Outside (PacketShaper 2500, 6500, 8500, 9500, and 10000), Right-Outside (PacketShaper 3500 and 7500), or Backup_Outside (PacketShaper 1400) 5 Lower_Inside (PacketShaper 2500, 6500, 8500, 9500, and 10000) or Left_Inside (PacketShaper 3500 and 7500) 6 Lower_Outside (PacketShaper 2500, 6500, 8500, 9500, and 10000) or Left_Outside (PacketShaper 3500 and 7500)
Packets	The total number of packets in the flow
Bytes	The total number of bytes in the flow
Start Time	Value of SysUpTime (the PacketShaper's up time) when the first packet in the flow was seen (in milliseconds)
End Time	Value of SysUpTime (the PacketShaper's up time) when the last packet in the flow was seen (in milliseconds)
Source Port	TCP or UDP port number of the device that the flow went out of
Destination Port	TCP or UDP port number of the device that the flow went to
Flow Policy ID	Type of PacketShaper policy assigned to the flow: 1 priority (shaping on) 2 rate (shaping on) 8 ignore (shaping on) 16 discard (shaping on) 32 never admit (shaping on) 48 server flow limited (shaping on) 49 client flow limited (shaping on) 129 priority (shaping off) 130 rate (shaping off) 136 ignore (shaping off) 144 discard (shaping off) 160 never admit (shaping off) 176 server flow limited (shaping off) 177 client flow limited (shaping off)
TCP Flag ID	Protocol state: (URG=32, ACK=16, PSH=8, RST=4, SYN=2, FIN=1)
Protocol ID	Type of layer 4 protocol. For example, ICMP=1, TCP=6, UDP=17.
ToS	Differentiated Services Code Point (DSCP) value that the flow was tagged with. This value designates special handling of traffic (precedence, delay, throughput, and reliability).
Service ID	Service type of the flow, represented as an identification number
Server Location	Location of the server for this flow: s source d destination 0 unknown The server location is useful for determining which of the source/destination hosts was acting as a server and which was the client (thus indicating which direction the

Column Name	Description
	bulk of the data should have flowed). Note that this may not be applicable to some protocols.
Priority ID	PacketShaper priority assigned to the flow (0-7); this could be the priority assigned with a priority policy or the priority assigned to excess rate with a rate policy
Retransmitted Bytes	Number of bytes that were retransmitted during the flow Retransmissions are a useful quality metric since each retransmission typically triggers a time-out delay on the client stack and thus a substantial delay in the application transaction.
VLAN ID	802.1q VLAN (Virtual LAN) identifier that the flow was tagged with
TTL	Time to Live — an upper limit on the number of routers through which a datagram may pass. With each router the datagram passes through, the TTL is decremented by one. When the field reaches 0, the datagram is thrown away and the sender is notified with an ICMP message. The purpose of this field is to limit packets from getting into an infinite routing loop.
Measure Type	<p>The value of Measure Type determines which response time measurement variables are recorded for the flow as follows:</p> <ul style="list-style-type: none"> p ping (ICMP) v RTCP (VoIP) a RTM t TCP 0 none <p>The following rows list the measurement variables that are recorded for each type of flow. The actual measurements (Measure 1, Measure 2, and Measure 3) are determined by the Measure Type value. The values for each type of response type measurement (ping, RTCP, RTM, or TCP) are described in the following sections.</p> <p>Note: The ftype_id was not an original part of the Packeteer-2 record format; it is used for internal DataCollector aggregation only.</p>
<p align="center">Response Time Measurement for ICMP (Measure Type=p)</p> <p><i>ICMP response time measurements are useful for tracking SLAs, such as server availability, that may be measured by network monitoring tools. These metrics apply only to ICMP flows.</i></p>	
Measure 1: Ping Delay	Average round trip delay between PacketShaper and destination
Measure 2: Ping Speed	Host speed of server in bps
Measure 3: Ping Success	Success percentage for pings
<p align="center">Response Time Measurement for VoIP (Measure type=v)</p> <p><i>Jitter, latency, and loss apply only to VoIP flows. These are end-to-end measurements derived from the internal feedback statistics inside the VoIP flow itself.</i></p>	

Column Name	Description
Measure 1: VoIP Latency	Packet exchange latency in milliseconds — PacketShaper to source and back to PacketShaper
Measure 2: VoIP Jitter	Average jitter on the RTP stream from destination to server
Measure 3: VoIP Loss	Percentage packet loss seen on the RTP stream from destination to source
Response Time Measurement for TCP (Measure type=a) <i>Server and network delay measurements help to indicate whether performance problems are coming from the server or from the network. These measurements are also useful as benchmarks to tell how much delay is being introduced into transactions and to compare how it deviates from reference points over time. These metrics apply only to TCP transaction-oriented traffic.</i>	
Measure 1: RTM Total Delay	Total transaction delay in ms
Measure 2: RTM Server Delay	Server delay all transactions, in ms
Measure 3: RTM Transactions	Total number of transactions in the flow
Packet Exchange Time (Measure Type=t) <i>Packet Exchange Time (PET) is a basic latency measurement for TCP traffic. Under normal conditions it should be an indication of the round trip time for packets from client to appliance and from appliance to server. This can be a useful replacement for ICMP since it is application specific, thus receiving the actual QoS of the transaction. Also it is measured on every real transaction. PET can also be useful for detecting times of congestion and estimating how much congestion was present.</i>	
Measure 1: Server Speed	Host speed of server in bps
Measure 2: PET Server	Average packet exchange time, PacketShaper to server and back
Measure 3: PET Client	Average packet exchange time, PacketShaper to client and back

Report Options

Each time you run an IntelligenceCenter report, you must set the report options that define the specific information—such as network group, sub-group, or view—for which to display data. This section lists the report options for each of the reports:

- [Application Activity Report Options](#)
- [Application Response Time Report Options](#)
- [Device Compression Report Options](#)
- [Host Pairs Activity Report Options](#)
- [Link Utilization Report Options](#)
- [Site Response Time Report Options](#)
- [TCP Health Report Options](#)
- [Top Applications Report Options](#)
- [Top Applications by Site Report Options](#)
- [Top DSCP Report Options](#)
- [Top Host Pairs Report Options](#)
- [Top Host Pairs by Site Report Options](#)
- [Top Immediate Children Report Options](#)
- [Top Listeners Report Options](#)
- [Top Listeners by Site Report Options](#)
- [Top Sites Report Options](#)
- [Top Services Report Options](#)
- [Top Services by Site Report Options](#)
- [Top Talkers Report Options](#)
- [Top Talkers by Site Report Options](#)
- [Top Traffic Classes Summary Report Options](#)
- [Top VLAN Report Options](#)
- [Traffic Class Utilization Report Options](#)
- [Traffic Class Compression Report Options](#)
- [Traffic Class Response Time Report Options](#)
- [VoIP Statistics Report Options](#)

Application Activity Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which to display application activity statistics. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Application	Click Select to display the Select Application Definition dialog box. You can then select the application to which to restrict the report and then click Accept .

Application Response Time Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which to display application response time statistics. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Critical Application	Specifies whether to restrict the report to applications that you have designated as critical .

Device Compression Summary Report Options

The Device Compression Summary report has a single report option: **Network Group**. This option allows you to select the portion of your network on which you want to view compression summary statistics. Click **Select** to display your [network](#) in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the **Devices** tab, or you can click the **View** tab to select a logical view to which to restrict the report. After you make your selection, click **Accept**.

Host Pairs Activity Report Options

Option	Description
Network Group or Device	Specifies the network group, sub-group, or view for which to display host pair activity. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Host IP Address	Specifies the IP address of the host that you are investigating. This host does not have to be a device within IC; usually it will be a client system on your network.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, Partition Size, Partition Utilization, Efficiency Level, Average Rate, Peak Rate.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Link Utilization Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which to display link utilization statistics. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Direction	Specify whether you want to display link utilization data for Inbound links, Outbound links, or Both inbound and outbound links.
Display Both Directions in	Specifies whether to display inbound and outbound link data on the Same chart or Separate charts . This field only displays if you set the Direction to Both .
Trending	Enable Trending if you want to include the Link Utilization Forecast on the report. You can use the forecast chart to help you estimate future link utilization based on historical data. You can only use trending if you select a time range of one day or more.
Forecast for next	Specifies the number of days, weeks, or months for which to graph link utilization forecast data, not to exceed a total of one year (365 days, 52 weeks, or 12 months). This field is only available if you enabled Trending .
Use actual data from last	Specifies the number of days, weeks, or months of historical data to use to calculate the link utilization forecast, not to exceed a total of one year (365 days, 52 weeks, or 12 months). This historical data will also be shown on the trending graph. This field is only available if you enabled Trending .
Ignore Outliers	Specifies whether to include values that deviate significantly from the rest of the data when calculating the link utilization forecast. If you enable this option, data values that are outside of the mean will be ignored when calculating the link utilization forecast. This field is only available if you enabled Trending .

Site Response Time Report Options

The Site Response Time report has a single report option: **Network Group**. This option allows you to select the portion of your network for which you want to view site statistics. Click **Select** to display your [network](#) in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the **Devices** tab, or you can click the **View** tab to select a logical view to which to restrict the report. After you make your selection, click **Accept**.

Site Response Time Report Options

The Site Response Time report has a single report option: **Network Group or Device**. This option allows you to select the portion of your network on which you want to view site statistics. Click **Select** to display your [network](#) in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the **Devices** tab, or you can click the **View** tab to select a logical view to which to restrict the report. After you make your selection, click **Accept**.

TCP Health Report Options

The TCP Health report has a single report option: **Network Group or Device**. This option allows you to select the portion of your network on which you want to view TCP health statistics. Click **Select** to display your [network](#) in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the **Devices** tab, or you can click the **View** tab to select a logical view to which to restrict the report. After you make your selection, click **Accept**.

Top Applications Report Options

Option	Description
Network Group or Device	This option allows you to select the portion of your network for which you want to view the top bandwidth-consuming applications. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Critical Application	Specifies whether to restrict the report to applications that you have designated as critical .
Max Count	Specifies how many top bandwidth-consuming applications to display on the report. For example, if you enter 10, the top 10 applications will be displayed; if you enter 15, the top 15 applications will be displayed.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Critical.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Applications by Site Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top applications by site. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Site	Allows you select the site for which you want to view application information. Click Select and then choose a site from the Select Site dialog box. You must have defined the site within IC before you can create reports based on the site.
Critical Application	Specifies whether to restrict the report to applications that you have designated as critical .
Max Count	Specifies how many top bandwidth-consuming applications to display on the report. For example, if you enter 10, the top 10 applications will be displayed; if you enter 15, the top 15 applications will be displayed.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Critical.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top DSCP Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top DSCP values. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Max Count	Specifies how many top DSCP values to display on the report, up to 200. For example, if you select 10, the top 10 DSCP values will be displayed; if you select 25, the top 25 DSCP values will be displayed. Keep in mind that the pie chart will only display the top 25 DSCP values, but any additional top DSCP values will be displayed on the corresponding table. Additionally, the higher the Max Count value, the longer the report will take to generate.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Flows.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Host Pairs Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top host pairs. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Max Count	Specifies how many top host pairs to display on the report, up to 200. For example, if you select 10, the top 10 host pairs will be displayed; if you select 25, the top 25 host pairs will be displayed. Keep in mind that the pie chart will only display the top 25 host pairs, but any additional top host pairs will be displayed on the corresponding table. Additionally, the higher the Max Count value, the longer the report will take to generate.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Flows. You can re-sort the report after you run it by clicking on the column heading.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Host Pairs by Site Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top host pairs by site. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Site	Allows you select the site for which you want to view host pair information. Click Select and then choose a site from the Select Site dialog box. You must have defined the site within IC before you can create reports based on the site.
Max Count	Specifies how many top bandwidth-consuming applications to display on the report. For example, if you enter 10, the top 10 host pairs will be displayed; if you enter 15, the top 15 host pairs will be displayed.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Flows.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Immediate Children Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top child traffic classes. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Traffic Class	Specifies the traffic class for which you want to display most active child classes. By default, the /Inbound traffic class is selected. To select a different parent class, click Select . The <i>Select Traffic Class</i> dialog box is displayed. Select a traffic class and then click Accept .
Max Count	Specifies how many top child classes to display on the report, up to 200. For example, if you enter 10, the top 10 children will be displayed; if you enter 25, the top 25 children will be displayed. Keep in mind that the pie chart will only display the top 25 immediate children, but any additional top children classes will be displayed on the corresponding table. Additionally, the higher the Max Count value, the longer the report will take to generate.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes or Packets.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Listeners by Site Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top listeners for the site. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Site	Allows you select the site for which you want to view top listeners. Click Select and then choose a site from the Select Site dialog box. You must have defined the site within IC before you can create reports based on the site.
Max Count	Specifies how many top listeners to display on the report. For example, if you enter 10, the top 10 listeners will be displayed; if you enter 15, the top 15 listeners will be displayed.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Flows.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Listeners Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top listeners. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Max Count	Specifies how many top traffic-receiving hosts to display on the report. For example, if you enter 10, the top 10 listeners will be displayed; if you enter 25, the top 25 listeners will be displayed. Keep in mind that the pie chart will only display the top 25 listeners, but any additional top listeners will be displayed on the corresponding table. Additionally, the higher the Max Count value, the longer the report will take to generate.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Flows.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Options for Top N by Site Reports

The following options apply to the following Top N site-based reports: Host Pairs, Listeners, Talkers, and Traffic Classes.

Option	Description
Network Group	Select the portion of your network on which you want to view application data. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Site	Select the name of the site for which you want to display the top N items.
Max Count	Specifies how many top items to display on the report. For example, in the Top N Applications by Site report, if you enter 10 for Max Count, the top 10 applications in the site will be displayed; if you enter 15, the top 15 applications will be displayed.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Flows.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Services by Site Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top services for a specific site. Click Select to display your network in a pop-up window. You can then select the network group, subgroup, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Site	The site for which to display top services. Click Select to display your sites in a pop-up window. Select the site to which to restrict the report and then click Accept .
Max Count	Specifies how many top traffic-generating hosts to display on the report, up to 200. For example, if you enter 10, the top 10 talkers will be displayed; if you enter 25, the top 25 talkers will be displayed. Keep in mind that the pie chart will only display the top 25 talkers, but any additional top talkers will be displayed on the corresponding table. Additionally, the higher the Max Count value, the longer the report will take to generate.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes (default), Packets, or Flows. You can re-sort the report after you run it by clicking on the column heading.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Services Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top services. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Max Count	Specifies how many top services to display on the report, up to 200. For example, if you enter 10, the top 10 services will be displayed; if you enter 25, the top 25 services will be displayed. Keep in mind that the pie chart will only display the top 25 services, but any additional top services will be displayed on the corresponding table. Additionally, the higher the Max Count value, the longer the report will take to generate.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Flows.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Sites Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top sites. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept . Note that you must have defined the site within IC before you can create site-based reports.
Max Count	Specifies how many top bandwidth-consuming sites to display on the report. For example, if you enter 10, the top 10 sites will be displayed; if you enter 15, the top 15 sites will be displayed.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes or Packets.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Talkers by Site Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top talkers for the site. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Site	Allows you select the site for which you want to view top talkers. Click Select and then choose a site from the Select Site dialog box. You must have defined the site within IC before you can create reports based on the site.
Max Count	Specifies how many top talkers to display on the report. For example, if you enter 10, the top 10 talkers will be displayed; if you enter 15, the top 15 talkers will be displayed.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Flows.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Talkers Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top talkers. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Max Count	Specifies how many top traffic-generating hosts to display on the report, up to 200. For example, if you enter 10, the top 10 talkers will be displayed; if you enter 25, the top 25 talkers will be displayed. Keep in mind that the pie chart will only display the top 25 talkers, but any additional top talkers will be displayed on the corresponding table. Additionally, the higher the Max Count value, the longer the report will take to generate.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Flows.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top Traffic Classes Summary Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top traffic classes. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Max Count	Specifies how many top traffic classes to display on the report, up to 200. For example, if you enter 10, the top 10 traffic classes will be displayed; if you enter 25, the top 25 traffic classes will be displayed. Keep in mind that the pie chart will only display the top 25 traffic classes, but any additional top traffic classes will be displayed on the corresponding table. Additionally, the higher the Max Count value, the longer the report will take to generate.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes or Packets.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Top VLAN Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view the top VLANs. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Max Count	Specifies how many top VLANs to display on the report, up to 200. For example, if you enter 10, the top 10 VLANs will be displayed; if you enter 15, the top 15 VLANs will be displayed. Keep in mind that the pie chart will only display the top 25 VLAN IDs, but any additional top VLAN IDs will be displayed on the corresponding table. Additionally, the higher the Max Count value, the longer the report will take to generate.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes, Packets, or Flows.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

Traffic Class Compression Report Options

The Traffic Class Compression report has a single report option: **Network Group or Device**. This option allows you to select the portion of your network on which you want to view compression statistics. Click **Select** to display your [network](#) in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the **Devices** tab, or you can click the **View** tab to select a logical view to which to restrict the report. After you make your selection, click **Accept**.

Traffic Class Response Time Report Options

The Traffic Class RTM report has a single report option: **Network Group or Device**. This option allows you to select the portion of your network on which you want to view traffic class RTM statistics. Click **Select** to display your [network](#) in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the **Devices** tab, or you can click the **View** tab to select a logical view to which to restrict the report. After you make your selection, click **Accept**.

Traffic Class Utilization Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which to display class utilization statistics. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Traffic Class	Allows you to select a traffic class on which to report. Click Select to display the Select Traffic Class dialog box. You can then select the traffic class to which to restrict the report and then click Accept .
Max Count	Specifies how many top traffic-generating child classes to display on the report, up to 200. For example, if you enter 10, the top 10 child classes will be displayed; if you enter 25, the top 25 child classes will be displayed. Keep in mind that the pie chart will only display the top 25 talkers, but any additional top talkers will be displayed on the corresponding table. Additionally, the higher the Max Count value, the longer the report will take to generate.
Sort Column	Specifies the return result column by which the report will be sorted by default. You can select from Bytes or Packets.
Sort Direction	Specifies whether the column on which you chose to sort the report will display in ascending or descending order.

VoIP Statistics Report Options

Option	Description
Network Group or Device	Specifies the portion of your network for which you want to view VoIP statistics. Click Select to display your network in a pop-up window. You can then select the network group, sub-group, or device to which to restrict the report on the Devices tab, or you can click the View tab to select a logical view to which to restrict the report. After you make your selection, click Accept .
Application	Specifies the VoIP application; you cannot select a different value. By default, the standard IC VoIP application includes all RTP-I, RTP-B, RTCP-I, and RTCP-B traffic. Keep in mind that you can customize the application so that it includes only the VoIP applications that you want to group. For best results, you should only group applications that use the same coder/decoder (codec) value.
Codec	Each VoIP application, device, or service uses a specific codec (coder/decoder) for the conversion between analog and digital signals. Select the codec used by the VoIP Application that you selected. To determine which codec to use, refer to the documentation provided by the VoIP application, device, or service provider. Selecting the correct codec ensures the most accurate reporting and calculation of VoIP metrics such as percent loss, latency, jitter, and MOS.

Index

A

Adding..... 26, 52, 75, 137, 164, 167
Applications ... 108, 110, 114, 115, 167, 170,
172, 202, 290
Archive.....124, 127, 132
Audit Log.....5

B

Backup.....21, 23, 43, 91

C

Compression.....206, 222
CSV Files.....43, 134, 316

D

Data granularity119
Database tables.....119
DataCollector . 23, 26, 28, 37, 43, 44, 46, 75,
82, 88, 96
Deleting.. 44, 56, 90, 93, 107, 115, 133, 172,
191
Device. 71, 75, 77, 80, 82, 84, 85, 87, 88, 90
Device Compression Summary206, 321
Disk Space.....23
DSCP.....240, 327

E

Email notification.....121, 124, 127

F

FDR28, 37, 43, 316

G

Granularity119

I

iShaper75, 88

L

Licensing.....7
Logging5, 20
Login3, 5, 12, 55, 87
Logout.....5

M

Metric data 25, 28, 37, 309, 314

Monitor 46, 96

N

NetFlow-5..... 25, 28, 37
Network branch..... 71, 73, 75
Network Device..... 75, 84, 88
Network view..... 102, 103, 105, 107
NTP 118

P

Packeteer-2..... 25, 28, 37, 316
PacketShaper 75, 77, 88, 91, 92, 93
Password 52, 54, 55
PolicyCenter..... 75, 80, 85
Portlet..... 162, 164, 191, 192, 299
 Application Performance Portlet. 167, 170,
 172, 290
 Per Server FDR Portlet 181, 184, 299
 Per Subnet FDR Portlet 25, 185, 301
 Portlet Views 162

R

Reports8, 120, 121, 124, 127, 129, 131, 132,
133, 134, 136, 137, 139, 193
 Application Activity 195, 321
 Application Response Time 197, 321
 Device Compression Summary... 206, 321
 Host Pairs Activity 238, 322
 TCP Health..... 216, 324
 Top N Applications 202, 325
 Top N DSCP 240, 327
 Top N Host Pairs 243, 328
 Top N Immediate Children 204, 330
 Top N Listeners 245, 332
 Top N Talkers..... 251, 337
 Top N Traffic Classes Summary . 220, 338
 Top N VLAN 253, 339
 Traffic Class Compression 222, 339
 VoIP Statistics 233, 341
Restore 22, 92

S

Scheduled Tasks	8, 136
Schedules ...	8, 127, 134, 136, 137, 139, 141
Services	19, 20
Single Sign-on	87
Sites	150, 151
SNTP	118
Subnet classes	259, 261, 267, 270

T

Topology	71, 73, 75
----------------	------------

Traffic Classes	36, 37, 108, 110, 114, 222
----------------------	----------------------------

Troubleshooting	20
-----------------------	----

U

Users.....	51, 52, 54, 55, 56, 87
------------	------------------------

V

Views	102, 103, 107
-------------	---------------

VoIP	233
------------	-----

W

Windows Services.....	19, 20
-----------------------	--------